Notice of References Cited	Application/Control No. 15/808,201	Applicant(s)/Patent Under Reexamination Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	Page 1 of 2

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-5142665-A	08-1992	Bigus; Joseph P.	G06N3/04	706/16
*	B	US-20010010034-A1	07-2001	Burton, John Mark	G06F17/5022	703/17
*	C	US-20060129506-A1	06-2006	Edelman; Gerald M.	G05D1/0088	706/12
*	D	US-7119810-B2	10-2006	Sumanaweera; Thilaka S.	G06T15/005	345/506
*	E	US-7219085-B2	05-2007	Buck; Ian Andrew	G06N3/08	706/12
*	F	US-7525547-B1	04-2009	Diard; Franck R.	G06T1/20	345/502
	G					
	H					
	I					
	J					
	K					
	L					
	M					

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	Cornwall et al., "Automatically Translating a General Purpose C++ Image Processing Library for GPUs", IEEE, 2006-06, 8 pages. (Year: 2006)
	V	Montrym et al., "The GeForce 6800", IEEE, 2005, 11 pages. (Year: 2005)
	W	Perumalla, Kalyan S., "Discrete-event Execution Alternatives on General Purpose Graphical Processing Units (GPGPUs)", IEEE, " Proceedings of the 20th Workshop on Principles of Advanced and Distributed Simulation (PADS'06)", 2006-03, 8 pages. (Year: 2006)
	X	Setoain et al., "Parallel Hyperspectral Image Processing on Commodity Graphics Hardware", IEEE, "Proceedings of the 2006 International Conference on Parallel Processing Workshops (ICPPW'06)", 2006-03, 8 pages. (Year: 2006)

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Notice of References Cited	Application/Control No. 15/808,201	Applicant(s)/Patent Under Reexamination Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	Page 2 of 2

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
	A					
	B					
	C					
	D					
	E					
	F					
	G					
	H					
	I					
	J					
	K					
	L					
	M					

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	Luo et al., "Artificial Neural Network Computation on Graphic Process Unit", IEEE, "Proceedings of International Joint Conference on Neural Networks, Montreal, Canada, July 31 - August 4, 2005", 2005-02, pp.622-626 (Year: 2005)
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/10/2018								
	1	✓								
	2	✓								
	3	✓								
	4	✓								
	5	✓								
	6	✓								
	7	✓								
	8	✓								
	9	✓								
	10	✓								
	11	✓								
	12	✓								
	13	✓								
	14	✓								
	15	✓								
	16	✓								
	17	✓								
	18	✓								
	19	✓								
	20	✓								
	21	✓								
	22	✓								
	23	✓								
	24	✓								
	25	✓								
	26	✓								
	27	✓								
	28	✓								
	29	✓								
	30	✓								
	31	✓								
	32	✓								
	33	✓								
	34	✓								
	35	✓								
	36	✓								
	37	✓								
	38	✓								
	39	✓								
	40	✓								
	41	✓								
	42	✓								

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018								
	43	✓								
	44	✓								
	45	✓								
	46	✓								
	47	✓								
	48	✓								
	49	✓								
	50	✓								
	51	✓								
	52	✓								
	53	✓								
	54	✓								
	55	✓								
	56	✓								
	57	✓								
	58	✓								
	59	✓								
	60	✓								
	61	✓								
	62	✓								
	63	✓								
	64	✓								
	65	✓								
	66	✓								
	67	✓								
	68	✓								
	69	✓								
	70	✓								
	71	✓								
	72	✓								
	73	✓								
	74	✓								
	75	✓								
	76	✓								
	77	✓								
	78	✓								
	79	✓								
	80	✓								
	81	✓								
	82	✓								
	83	✓								
	84	✓								
	85	✓								
	86	✓								
	87	✓								
	88	✓								
	89	✓								
	90	✓								
	91	✓								
	92	✓								
	93	✓								
	94	✓								

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018								
	95	✓								
	96	✓								
	97	✓								
	98	✓								
	99	✓								
	100	✓								
	101	✓								
	102	✓								
	103	✓								
	104	✓								
	105	✓								
	106	✓								
	107	✓								
	108	✓								
	109	✓								
	110	✓								
	111	✓								
	112	✓								
	113	✓								
	114	✓								
	115	✓								
	116	✓								
	117	✓								
	118	✓								
	119	✓								
	120	✓								
	121	✓								
	122	✓								
	123	✓								
	124	✓								
	125	✓								
	126	✓								
	127	✓								
	128	✓								
	129	✓								
	130	✓								
	131	✓								
	132	✓								
	133	✓								
	134	✓								
	135	✓								
	136	✓								

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

CPC - Searched*		
Symbol	Date	Examiner
G06F 9/5027	10/10/2018	/WHW/
G06T 1/20	10/10/2018	/WHW/
G06T 1/60	10/10/2018	/WHW/
G06N 99/005	10/10/2018	/WHW/
G06N 3/063	10/10/2018	/WHW/
G06F 2209/509	10/10/2018	/WHW/

CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

Search Notes		
Search Notes	Date	Examiner
Prosecution History Review	10/10/2018	/WHW/
IDS Review	10/10/2018	/WHW/
Litigation Search	10/10/2018	/WHW/
EAST BRS Keyword Search	10/10/2018	/WHW/
NPL IEEE Search	10/10/2018	/WHW/

/William H. Wood/ Primary Examiner, Art Unit 3992	

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

Interference Search			
US Class/CPC Symbol	US Subclass/CPC Group	Date	Examiner

/William H. Wood/ Primary Examiner, Art Unit 3992	

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/10/09 15:56
L2	1	L1 and pointer	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 15:56
S1	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/06/15 17:23
S2	12	("20020064314" | "20020168100" | "20030026588" | "20090080695" | "20090116688" | "20110279682" | "20130126703" | "20130131985" | "20150224648" | "5063603" | "6018696" | "8583286").PN.	US-PGPUB; USPAT	OR	ON	2018/06/22 16:39
S3	47	("20020046271" | "20020050518" | "20030078754" | "20040015334" | "20050166042" | "20060184273" | "20070052713" | "20070198222" | "20070279429" | "20080033897" | "20080066065" | "20080258880" | "20090089030" | "20100048242" | "20110004341" | "20110173015" | "20120072215" | "20120089552" | "20120197596" | "20120316786" | "20140019392" | "20140032461" | "20140052679" | "20140089232" | "20150127149" | "20150134232" | "20160075017" | "20160082597" | "20160096270" | "20160198000" | "20170024877" | "20170076194" | "20170193298" | "5136687" | "5172253" | "5388206" | "6336051" | "6647508" | "7765029" | "7861060" | "7873650" | "8392346" | "8510244" | "8648867" | "9031692" | "9177246" | "9626566").PN.	US-PGPUB; USPAT	OR	ON	2018/06/22 16:48
S4	1	("8648867").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/09/07 11:39
S5	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/09/07 11:44
S6	1	S5 and neuron	US-PGPUB; USPAT; USOCR;	OR	ON	2018/09/07 11:44

			FPRS; EPO; JPO; DERWENT; IBM_TDB			
S7	4	"20140192073"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/11 15:53
S9	1	("20080117220").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/09/11 16:01
S10	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/09/11 16:35
S11	1	S10 and cycle	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/11 16:35
S12	1	S10 and pointer	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/11 16:37
S13	1	S10 and intermediate	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/11 16:38
S14	4591	G06F9/5027	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:13
S15	26814	G06T1/20	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:13
S16	19812	G06T1/60	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:13
S17	15204	G06N99/005	US-PGPUB;	OR	ON	2018/09/12

			USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			14:14
S18	5688	G06N3/063	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:14
S19	1273	G06F2209/509	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:14
S20	0	anatoli.in. adj gorchetchnikov.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:15
S21	0	(anatoli adj gorchetchnikov).in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:15
S22	0	(anatoli adj gorchetchnikov)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:16
S23	0	anatoli-gorchetchnikov.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:17
S24	0	heather-marie-ames.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:18
S25	0	heather-ames.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT;	OR	ON	2018/09/12 14:18

			IBM_TDB			
S26	0	massimiliano-versace.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:18
S27	0	fabrizio-santini.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:19
S28	7	anatoli.in. and gorchetchnikov.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:19
S29	15	heather.in. and ames.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:20
S30	28	massimiliano.in. and versace.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:21
S31	23	fabrizio.in. and santini.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:22
S32	1	(S14 or S15 or S16 or S17 or S18 or S19) and (neural with simulation with (graphic$4 adj processing adj unit))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:35
S33	53	(S14 or S15 or S16 or S17 or S18 or S19) and (simulation with (graphic$4 adj processing adj unit))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 14:37
S34	1432	(S14 or S15 or S16 or S17 or S18 or S19) and (artificial adj neural adj network)	US-PGPUB; USPAT; USOCR; FPRS;	OR	ON	2018/09/12 15:12

			EPO; JPO; DERWENT; IBM_TDB			
S35	232	(S14 or S15 or S16 or S17 or S18 or S19) and (artificial adj neural adj network) and (GPU or (graphic$4 adj processing adj unit))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:13
S36	16	(S14 or S15 or S16 or S17 or S18 or S19) and (artificial adj neural adj network) and (GPU or (graphic$4 adj processing adj unit)) and (accumulator)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:14
S37	21	(S14 or S15 or S16 or S17 or S18 or S19) and (artificial adj neural adj network) and (GPU or (graphic$4 adj processing adj unit)) and (summation)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:15
S38	0	(S14 or S15 or S16 or S17 or S18 or S19) and ((GPU or (graphic$4 adj processing adj unit)) near3 (output$4 or result$4) near3 (neural or neuron))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:25
S39	4	(S14 or S15 or S16 or S17 or S18 or S19) and ((GPU or (graphic$4 adj processing adj unit)) with (output$4 or result$4) with (neural or neuron))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:26
S40	35	(S14 or S15 or S16 or S17 or S18 or S19) and ((GPU or (graphic$4 adj processing adj unit)) and ((multip$4 or plural$4 or several or many) near3 memory) and ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:31
S41	1	(S14 or S15 or S16 or S17 or S18 or S19) and ((GPU or (graphic$4 adj processing adj unit)) and (((multip$4 or plural$4 or several or many) near3 memory) with ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4)))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:32
S42	1	(S14 or S15 or S16 or S17 or S18 or S19) and ((GPU or (graphic$4 adj processing adj unit)) and (((multip$4 or plural$4 or several or many) near3 memory) same ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4)))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:32
S43	26	(S14 or S15 or S16 or S17 or S18 or S19) and (((multip$4 or plural$4 or several or	US-PGPUB; USPAT;	OR	ON	2018/09/12 15:38

		many) near3 memory) with ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4))	USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			
S44	0	(S14 or S15 or S16 or S17 or S18 or S19) and (((multip$4 or plural$4 or several or many) near3 (memory adj partition)) with ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:39
S45	0	(S14 or S15 or S16 or S17 or S18 or S19) and (((multip$4 or plural$4 or several or many) near3 (memory adj partition)) same ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:39
S46	0	(outsourc$4 or (handoff$4 or hand-off$4 or (hand adj off$4))) near3 simulation near3 ((graphic$4 adj processing adj unit) or GPU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:51
S47	1	(outsourc$4 or (handoff$4 or hand-off$4 or (hand adj off$4))) with simulation with ((graphic$4 adj processing adj unit) or GPU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:52
S48	0	(outsourc$4 or (handoff$4 or hand-off$4 or (hand adj off$4))) with (neural or neuron) with ((graphic$4 adj processing adj unit) or GPU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 15:53
S49	1	(outsourc$4 or (handoff$4 or hand-off$4 or (hand adj off$4))) same simulation same ((graphic$4 adj processing adj unit) or GPU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:46
S50	0	(outsourc$4 or (handoff$4 or hand-off$4 or (hand adj off$4))) same (neural or neuron) same ((graphic$4 adj processing adj unit) or GPU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:46
S51	0	(((multip$4 or plural$4 or several or many) near3 (memory adj partition)) with ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:47

S52	0	(((multip$4 or plural$4 or several or many) near3 (memory adj partition)) same ((parallel or concurrent$4 or simultaneous$4) near3 transfer$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:48
S53	584	(multip$4 or plural$4 or several or many) near3 (memory adj partition)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:48
S54	12747	(multip$4 or plural$4 or several or many) near3 (memory adj (bank or partition or section))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:49
S55	30053	(multip$4 or plural$4 or several or many) near3 (memory adj (bank or partition or section or module))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:50
S56	65	(multip$4 or plural$4 or several or many) near3 (memory adj (bank or partition or section or module)) with ((graphic$4 adj processing adj unit) or GPU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/12 16:51
S57	12	("20020064314" | "20020168100" | "20030026588" | "20090080695" | "20090116688" | "20110279682" | "20130126703" | "20130131985" | "20150224648" | "5063603" | "6018696" | "8583286").PN.	US-PGPUB; USPAT	OR	ON	2018/09/20 12:54
S58	47	("20020046271" | "20020050518" | "20030078754" | "20040015334" | "20050166042" | "20060184273" | "20070052713" | "20070198222" | "20070279429" | "20080033897" | "20080066065" | "20080258880" | "20090089030" | "20100048242" | "20110004341" | "20110173015" | "20120072215" | "20120089552" | "20120197596" | "20120316786" | "20140019392" | "20140032461" | "20140052679" | "20140089232" | "20150127149" | "20150134232" | "20160075017" | "20160082597" | "20160096270" | "20160198000" | "20170024877" | "20170076194" | "20170193298" | "5136687" | "5172253" | "5388206" | "6336051" | "6647508" | "7765029" | "7861060" | "7873650" | "8392346" | "8510244" | "8648867" | "9031692" | "9177246" |	US-PGPUB; USPAT	OR	ON	2018/09/20 12:54

NVIDIA Ex 1002

		"9626566").PN.				
S59	6	(S57 or S58) and swap$4	US-PGPUB; USPAT	OR	ON	2018/09/20 12:54
S60	1479	(general adj purpose adj computing) near2 (graphics adj processing adj unit)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 13:47
S61	1055	S60 and simulation	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 13:48
S62	48	S61 and (neuron or neural)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 13:49
S63	6	S61 and ((parallel or simultaneous$4 or concurrent$4) near3 (transfer$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 13:56
S64	53	S61 and ((parallel or simultaneous$4 or concurrent$4) near3 (memory))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 13:57
S65	14	S61 and ((parallel or simultaneous$4 or concurrent$4) near3 (input))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 13:59
S66	14	S61 and ((parallel or simultaneous$4 or concurrent$4) near3 (input$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 14:00
S67	19	S61 and ((parallel or simultaneous$4 or concurrent$4) near3 (output$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 14:00
S68	1364	S60 and (parallel or simultaneous$4 or	US-PGPUB;	OR	ON	2018/09/20

		concurrent$4)	USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			14:01
S71	2	S68 and "2006".ay.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 14:05
S72	1015	S68 and accelerator	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 14:11
S73	88	S68 and (accelerator with memory)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 14:11
S74	13	S68 and (dependent$4 with thread$4)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 14:16
S75	207	(accelerator adj memory) and (graphics adj processing adj unit)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 16:07
S76	41750	(graphics adj processing adj unit) with memory	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 16:29
S77	422	(graphics adj processing adj unit) with memory with transfer$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/20 16:30
S78	37	(graphics adj processing adj unit) with memory with transfer$4 with (parallel or simultaneous$4 or concurrent$4)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT;	OR	ON	2018/09/20 16:30

			IBM_TDB			
S79	82	(graphics adj processing adj unit) with swap$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:26
S80	397	(GPU) with swap$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:42
S81	6	(GPU) with swap$4 with pointer	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:42
S82	0	(GPGPU) with swap$4 with pointer	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:47
S83	22	(GPGPU) and (swap$4 with pointer)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:47
S84	3490	(GPGPU) and (during)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:49
S85	1070	(GPGPU) and (during adj (perform$4 or operat$4 or process$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:50
S86	135	(GPGPU) and ((during adj (perform$4 or operat$4 or process$4)) with memory)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 09:51
S87	1	("9189828").URPN.	USPAT	OR	ON	2018/09/21 09:52
S88	24	("20020046271" | "20030078754" |	US-PGPUB;	OR	ON	2018/09/21

		"20040015334" | "20050166042" | "20070052713" | "20070198222" | "20070279429" | "20080033897" | "20080117220" | "20080258880" | "20090089030" | "20100048242" | "20110173015" | "20120316786" | "20140192073" | "20160082597" | "20160198000" | "5388206" | "6336051" | "6647508" | "7861060" | "7873650" | "8648867" | "9189828").PN.	USPAT; USOCR			09:52
S89	3	("2008/0117220").URPN.	USPAT	OR	ON	2018/09/21 09:56
S90	314	(video near3 memory) same (GPU or (graphics adj processing adj unit)) same (transfer$4 near3 data)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 10:20
S91	55	(video near3 memory) same (GPU or (graphics adj processing adj unit)) same (transfer$4 near3 data) same (main adj memory)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 10:21
S92	2404	(GPU or (graphics adj processing adj unit)) with transfer$4 with (memory or buffer or cache)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 11:03
S93	264	(GPU or (graphics adj processing adj unit)) with transfer$4 with (memory or buffer or cache) with (parallel or simultaneous$4 or concurrent$4)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/21 11:04
S94	3723	GPU and ((parallel or simultaneous$4 or concurrent$4 or overlap$4) near3 (read$4 or writ$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/25 16:25
S95	287	GPU with ((parallel or simultaneous$4 or concurrent$4 or overlap$4) near3 (read$4 or writ$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/25 16:25
S96	1	("7219085").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/09/26 13:18
S97	2	(("7477256") or ("8212831")).PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/09/26 14:06

S98	1	("7219085").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/09/28 08:33
S99	0	S98 and simul$6	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/28 08:33
S100	0	S98 and pointer	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/28 08:51
S101	1	S98 and swap$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/09/28 08:51
S102	34	(GPU or (graphical adj processing adj unit)) same (numerical adj simulation)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/01 12:19
S103	1	(numerical adj simulation) and (swap$4 with intermediate)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/02 14:00
S104	2	((GPU or (graphical adj processing adj unit)) same (simulation)) and (swap$4 with intermediate)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/02 14:01
S105	153	(simulation) and (swap$4 with intermediate)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/02 14:03
S106	348	(simulation) and (swap$4 with pointer)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/02 14:03
S107	41	(US-20170102950-$ or US-20070124363-$ or US-20030152288-$ or	US-PGPUB; USPAT;	OR	ON	2018/10/03 15:07

		US-20080238917-$ or US-20080001957-$ or US-20040207618-$ or US-20020190994-$ or US-20070052713-$ or US-20060106591-$ or US-20110035737-$ or US-20090222572-$ or US-20090027402-$ or US-20090027383-$ or US-20080181081-$ or US-20060233525-$ or US-20060114254-$ or US-20010052123-$ or US-20070165035-$ or US-20010010034-$).did. or (US-7119810-$ or US-7746350-$ or US-9349201-$ or US-8194085-$ or US-7737984-$ or US-7068278-$ or US-7522167-$ or US-6999603-$ or US-7526634-$ or US-7525547-$ or US-6367122-$ or US-6146277-$ or US-7219085-$ or US-7477256-$ or US-6952827-$ or US-5977977-$ or US-5099447-$).did. or (WO-9316438-$ or WO-0227658-$).did. or (US-20050125369-$ or US-20050128206-$).did.	FPRS; DERWENT			
S108	10	S107 and (stor$4 near3 disk) and (parallel$4 or simultaneous$4 or concurrent$4)	US-PGPUB; USPAT; FPRS; DERWENT	OR	ON	2018/10/03 15:08
S109	0	(sav$4 adj "to" adj disk)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/04 15:07
S110	17260	(sav$4 near2 disk)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/04 15:07
S111	126	S110 with simulation	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/04 15:08
S112	2	S111 and GPU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/04 15:08
S113	1	("7219085").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/10/05 12:10
S114	1	S113 and logical	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO;	OR	ON	2018/10/05 12:10

			DERWENT; IBM_TDB			
S115	43	(US-20170102950-$ or US-20070124363-$ or US-20030152288-$ or US-20080238917-$ or US-20080001957-$ or US-20040207618-$ or US-20020190994-$ or US-20070052713-$ or US-20060106591-$ or US-20110035737-$ or US-20090222572-$ or US-20090027402-$ or US-20090027383-$ or US-20080181081-$ or US-20060233525-$ or US-20060114254-$ or US-20010052123-$ or US-20070165035-$ or US-20010010034-$ or US-20060129506-$).did. or (US-7119810-$ or US-7746350-$ or US-9349201-$ or US-8194085-$ or US-7737984-$ or US-7068278-$ or US-7522167-$ or US-6999603-$ or US-7526634-$ or US-7525547-$ or US-6367122-$ or US-6146277-$ or US-7219085-$ or US-7477256-$ or US-6952827-$ or US-5977977-$ or US-5099447-$ or US-5987443-$).did. or (WO-9316438-$ or WO-0227658-$).did. or (US-20050125369-$ or US-20050128206-$).did.	US-PGPUB; USPAT; FPRS; DERWENT	OR	ON	2018/10/05 12:11
S116	14	S115 and logical	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/05 12:11
S117	18	S115 and shared	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/05 12:14
S118	1	("20010010034").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2018/10/09 14:39
S119	0	S118 and display$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 14:39
S120	20725	display$4 near3 simulation	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 14:40
S121	934	video near3 display$4 near3 simulation	US-PGPUB; USPAT; USOCR; FPRS;	OR	ON	2018/10/09 14:40

			EPO; JPO; DERWENT; IBM_TDB			
S122	0	S121 and UML	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 14:41
S123	38	S121 and GPU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 14:41
S124	2546	neur$4 near3 simulation	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 14:43
S125	1269	S124 and display$4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 14:43
S126	153	S124 and (display$4 near3 result)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2018/10/09 14:44

10/9/2018 4:13:41 PM
C:\Users\wwood\Documents\EAST\Workspaces\15808201.wsp

PTO/SB/08a (07-09)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	3992
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code2 (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
	1.	5,063,603	11-05-1991	Burt	
	2.	6,018,696	01-25-2000	Matsuoka et al.	
	3.	8,583,286	11-12-2013	Fleischer et al.	
	4.	2002/0064314	05-30-2002	Comaniciu et al.	
	5.	2002/0168100	11-14-2002	Woodall	
	6.	2003/0026588	02-06-2003	Elder et al.	
	7.	2009/0080695	03-26-2009	Yang et al.	
	8.	2009/0116688	05-07-2009	Monacos et al.	
	9.	2011/0279682	11-17-2011	Li et al.	
	10.	2013/0126703	05-23-2013	Caulfield	
	11.	2013/0131985	05-23-2013	Weiland et al.	
	12.	2015/0224648	08-13-2015	Lee et al.	

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
	13.	Non-Final Office Action mailed 01/04/2018 from U.S. Application Serial No. 15/262,637, 23 pages (Atty. Docket No. NRLA-004/01US)	
	14.	WU, YAN & J. CAI, H. (2010). A Simulation Study of Deep Belief Network Combined with the Self-Organizing Mechanism of Adaptive Resonance Theory. 10.1109/CISE.2010.5677265, 4 pages.	

167642600 v1

Examiner Signature:	/William H. Wood/	Date Considered	06/22/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

PTO/SB/08a (07-09)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT (Use as many sheets as necessary)	Complete if Known	
	Application Number	15/808,201
	Filing Date	11-09-2017
	First Named Inventor	Anatoli Gorchetchnikov
	Art Unit	3992
	Examiner Name	WOOD, WILLIAM H.
	Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
	1.	AL-KAYSI, A. M. et al., A Multichannel Deep Belief Network for the Classification of EEG Data, from Ontology-based Information Extraction for Residential Land Use Suitability: A Case Study of the City of Regina, Canada, DOI 10.1007/978-3-319-26561-2_5, 8 pages (November 2015).	
	2.	APOLLONI, B. et al., Training a network of mobile neurons, Proceedings of International Joint Conference on Neural Networks, San Jose, CA, doi: 10.1109/IJCNN.2011.6033427, pp. 1683-1691 (July 31-August 5, 2011).	
	3.	BODDAPATI, V., Classifying Environmental Sounds with Image Networks, Thesis, Faculty of Computing Blekinge Institute of Technology, 37 pages (February 2017).	
	4.	KHALIGH-RAZAVI, S. -M. et al., Deep Supervised, but Not Unsupervised, Models May Explain IT Cortical Representation, PLoS Computational Biology, Vol. 10, Issue 11, 29 pages (November 2014).	
	5.	KIM, S., Novel approaches to clustering, biclustering and algorithms based on adaptive resonance theory and intelligent control, Doctoral Dissertations, Missouri University of Science and Technology, 125 pages (2016).	
	6.	Notice of Alllowance mailed May 22, 2018 from U.S. Application Serial No. 15/262,637, 6 pages (Atty. Docket No. NRLA-004/01US).	
	7.	REN, Y. et al., Ensemble Classification and Regression - Recent Developments, Applications and Future Directions, in IEEE Computational Intelligence Magazine, 10.1109/MCI.2015.2471235, 14 pages (2016).	
	8.	SUN, Z. et al., Recognition of SAR target based on multilayer auto-encoder and SNN, International Journal of Innovative Computing, Information and Control, Vol. 9, No. 11, pp. 4331-4341, November 2013.	

177153193 v1

Examiner Signature:	/William H. Wood/	Date Considered	06/22/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.
[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within results | Alert! Download PDFs | Export | Set Search Alerts | Search History!

Displaying 1 of 1 result for ("neuron" NEAR/6 "GPU")

Filters Applied: 2006

☐ Journals & Magazines (1)

Visualization of Fibrous and Thread-like Data

Zeki Melek | David Mayerich | Cem Yuksel | John Keyser
IEEE Transactions on Visualization and Computer Graphics
Year: 2006 | Volume: 12 | Issue: 5

☐ **Visualization of Fibrous and Thread-like Data**
Zeki Melek ; David Mayerich ; Cem Yuksel ; John Keyser
IEEE Transactions on Visualization and Computer Graphics
Year: 2006, Volume: 12, Issue: 5
Pages: 1165 - 1172
Cited by: Papers (14) | Patents (1)
IEEE Journals & Magazines

Abstract (5975 Kb)

« First < 1 > Next >>

IEEE Account
» Change Username/Password
» Update Address

Purchase Details
» Payment Options
» Order History
» View Purchased Documents

Profile Information
» Communications Preferences
» Profession and Education
» Technical Interests

Need Help?
» **US & Canada:** +1 800 678 4333
» **Worldwide:** +1 732 981 0060
» Contact & Support

| About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

IEEE Account

Profile Information

Purchase Details

Need Help?

Other

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

US & Canada: +1 800 678 4333
Worldwide: +1 732 981 0060

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within result: Filter | Alerts Show: Show All Results Download PDFs Export | Set Search Alerts Search History Sort By: Sort Relevance

Displaying results 1-2 of 2 for ("neural" NEAR/6 "GPU")

Filters Applied: 2005

☐ Conferences (2)

Filter results by

Search within results

Content Type

☐ Conferences (2)

Year

Single Year | Range

2005 2005

From To

2005 2005

Author

Affiliation

Publication Title

Publisher

Conference Location

Artificial neural network computation on graphic process unit

Zhongwen Luo | Hongzhi Liu | Xincai Wu
Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.
Year: 2005 | Volume: 1

☐ **Artificial neural network computation on graphic process unit**
Zhongwen Luo; Hongzhi Liu; Xincai Wu
Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.
Year: 2005, Volume: 1
Pages: 622 - 626 vol. 1
Cited by: Papers (8)
IEEE Conferences

Abstract (2239 Kb)

Using GPUs for machine learning algorithms

D. Steinkraus | I. Buck | P.Y. Simard
Eighth International Conference on Document Analysis and Recognition (ICDAR'05)
Year: 2005

☐ **Using GPUs for machine learning algorithms**
D. Steinkraus ; I. Buck ; P.Y. Simard
Eighth International Conference on Document Analysis and Recognition (ICDAR'05)
Year: 2005
Pages: 1115 - 1120 Vol. 2
Cited by: Papers (29) | Patents (8)
IEEE Conferences

Abstract (4349 Kb)



« First < 1 > Next >>

IEEE Account
» Change Username/Password
» Update Address

Purchase Details
» Payment Options
» Order History
» View Purchased Documents

Profile Information
» Communications Preferences
» Profession and Education
» Technical Interests

Need Help?
» US & Canada: +1 800 678 4333
» Worldwide: +1 732 981 0060
» Contact & Support

| About IEEE Xplore | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

A not-for-profit organization, IEEE is the world's largest professional association dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

IEEE Account

Profile Information

Purchase Details

Need Help?

Other

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

US & Canada: +1 800 678 4333
Worldwide: +1 732 981 0060

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within results | Alert! Download PDFs | Export | Set Search Alerts | Search History!

Displaying 1 of 1 result for ("neural" AND "GPGPU")

Filters Applied: 2007

☐ Conferences (1)

Approximate Dynamic Programming and Neural Networks on Game Hardware

Ryan J. Meuth | Donald C. Wunsch
2007 International Joint Conference on Neural Networks
Year: 2007

☐ **Approximate Dynamic Programming and Neural Networks on Game Hardware**
Ryan J. Meuth | Donald C. Wunsch
2007 International Joint Conference on Neural Networks
Year: 2007
Pages: 852 - 856
Cited by: Papers (3)
IEEE Conferences

Abstract (250 Kb)

« First < 1 > Next >>

IEEE Account
» Change Username/Password
» Update Address

Purchase Details
» Payment Options
» Order History
» View Purchased Documents

Profile Information
» Communications Preferences
» Profession and Education
» Technical Interests

Need Help?
» **US & Canada:** +1 800 678 4333
» **Worldwide:** +1 732 981 0060
» Contact & Support

| About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

IEEE Account

Profile Information

Purchase Details

Need Help?

Other

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.

US & Canada: +1 800 678 4333
Worldwide: +1 732 981 0060

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within results | Alert! Download PDFs | Export | Set Search Alerts | Search History!

Displaying 1 of 1 result for ("neuron" AND "GPGPU")

Filters Applied: 2010

☐ Conferences (1)

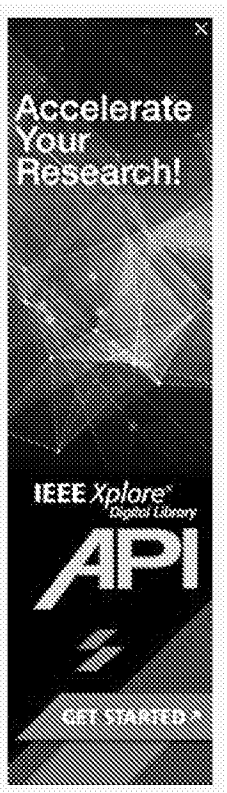
Neuromorphic models on a GPGPU cluster

Bing Han | Tarek M. Taha
The 2010 International Joint Conference on Neural Networks (IJCNN)
Year: 2010

☐ **Neuromorphic models on a GPGPU cluster**
Bing Han ; Tarek M. Taha
The 2010 International Joint Conference on Neural Networks (IJCNN)
Year: 2010
Pages: 1 - 8
Cited by: Papers (2)
IEEE Conferences

Abstract (1249 Kb)

« First < 1 > Next >>



IEEE Account
» Change Username/Password
» Update Address

Purchase Details
» Payment Options
» Order History
» View Purchased Documents

Profile Information
» Communications Preferences
» Profession and Education
» Technical Interests

Need Help?
» US & Canada: +1 800 678 4333
» Worldwide: +1 732 981 0060
» Contact & Support

| About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

IEEE Account
Profile Information
Purchase Details

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

US & Canada: +1 800 678 4333
Worldwide: +1 732 981 0060

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within result: Filter | Alert Show: Show All Results Download PDFs Export | Set Search Alerts Search History Sort By: Sort Relevance

Displaying results 1-2 of 2 for ("simulation" AND "GPGPU")

Filters Applied: 2005-2006

☐ Conferences (2)

Filter results by

Search within results

Content Type

☐ Conferences (2)

Year

Single Year | Range

☐ 2005 ☐ 2006

From 2005 To 2006

Author

Affiliation

Publication Title

Publisher

Conference Location

Discrete-event Execution Alternatives on General Purpose Graphical Processing Units (GPGPUs)

K.S. Perumalla
20th Workshop on Principles of Advanced and Distributed Simulation (PADS'06)
Year: 2006

☐ **Discrete-event Execution Alternatives on General Purpose Graphical Processing Units (GPGPUs)**
K.S. Perumalla
20th Workshop on Principles of Advanced and Distributed Simulation (PADS'06)
Year: 2006
Pages: 74 - 81
Cited by: Papers (27)
IEEE Conferences

Abstract (289 Kb)

VolQD: direct volume rendering of multi-million atom quantum dot simulations

W. Qiao | D.S. Ebert | A. Entezari | M. Korkusinski | G. Klimeck
VIS 05. IEEE Visualization, 2005.
Year: 2005

☐ **VolQD: direct volume rendering of multi-million atom quantum dot simulations**
W. Qiao ; D.S. Ebert ; A. Entezari ; M. Korkusinski ; G. Klimeck
VIS 05. IEEE Visualization, 2005.
Year: 2005
Pages: 319 - 326
Cited by: Papers (5)
IEEE Conferences

Abstract (1779 Kb)

« First < 1 > Next >>

IEEE Account
» Change Username/Password
» Update Address

Purchase Details
» Payment Options
» Order History
» View Purchased Documents

Profile Information
» Communications Preferences
» Profession and Education
» Technical Interests

Need Help?
» US & Canada: +1 800 678 4333
» Worldwide: +1 732 981 0060
» Contact & Support

| About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.

© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

IEEE Account

Profile Information

Purchase Details

Need Help?

Other

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

US & Canada: +1 800 678 4333
Worldwide: +1 732 981 0060

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within result: Filter | Alert Show: Show All Results Download PDFs Export | Set Search Alerts Search History Sort By: Sort Relevance

Displaying results 1-2 of 2 for ("GPU" and ("main memory" NEAR/3 "output"))

Filters Applied: 1974-2006

☐ Conferences (1) ☐ Journals & Magazines (1)

Filter results by

Search within results

Content Type

☐ Conferences (1)

☐ Journals & Magazines (1)

Year

Single Year | Range

☐ 1974 From

☐ 2006 To

1974

2006

Author

Affiliation

Publication Title

Publisher

Conference Location

Automatically translating a general purpose C++ image processing library for GPUs

J.L.T. Cornwall | O. Beckmann | P.H.J. Kelly
Proceedings 20th IEEE International Parallel & Distributed Processing Symposium
Year: 2006

☐ **Automatically translating a general purpose C++ image processing library for GPUs**
J.L.T. Cornwall ; O. Beckmann ; P.H.J. Kelly
Proceedings 20th IEEE International Parallel & Distributed Processing Symposium
Year: 2006
Page: 8 pp
Cited by: Papers (1)
IEEE Conferences

Abstract (229 Kb)

Microprogramming: Perspective and Status

A.K. Agrawala | T.G. Rauscher
IEEE Transactions on Computers
Year: 1974 | Volume: C-23 | Issue: 8

☐ **Microprogramming: Perspective and Status**
A.K. Agrawala ; T.G. Rauscher
IEEE Transactions on Computers
Year: 1974, Volume: C-23, Issue: 8
Pages: 817 - 837
Cited by: Papers (22) | Patents (2)
IEEE Journals & Magazines

Abstract (6428 Kb)

« First < 1 > Next >>

IEEE Account
» Change Username/Password
» Update Address

Purchase Details
» Payment Options
» Order History
» View Purchased Documents

Profile Information
» Communications Preferences
» Profession and Education
» Technical Interests

Need Help?
» US & Canada: +1 800 678 4333
» Worldwide: +1 732 981 0060
» Contact & Support

| About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

A public charity, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

IEEE Account

Profile Information

Purchase Details

Need Help?

Other

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

US & Canada: +1 800 678 4333
Worldwide: +1 732 981 0060

IEEE websites place cookies on your device to give you the best user experience. By using our websites, you agree to the placement of these cookies. To learn more, read our Privacy Policy.

Accept & Close

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within result: Filter | Alert Show: Show All Results Download PDFs Export | Set Search Alerts Search History Sort By: Sort Relevance

Displaying results 1-7 of 7 for ("GPU" and ("main memory" NEAR/3 "transfer"))

Filters Applied: 2004-2006

☐ Conferences (4) ☐ Journals & Magazines (3)

Filter results by

Search within results

Content Type

☐ Conferences (4)

☐ Journals & Magazines (3)

Year

Single Year | Range

☐ 2004 ☐ 2006

From 2004 To 2006

Author

Affiliation

Publication Title

Publisher

Conference Location

Expressway to Your Skull

E. Guizzo | H. Goldstein
IEEE Spectrum
Year: 2006 | Volume: 43 | Issue: 8

☐ **Expressway to Your Skull**
E. Guizzo ; H. Goldstein
IEEE Spectrum
Year: 2006, Volume: 43, Issue: 8
Pages: 34 - 39
Cited by: Papers (1)
IEEE Journals & Magazines

Abstract (10136 Kb)

Real-time multi-stereo depth estimation on GPU with approximative discontinuity handling

J. Woetzel | R. Koch
1st European Conference on Visual Media Production (CVMP) 2004
Year: 2004

☐ **Real-time multi-stereo depth estimation on GPU with approximative discontinuity handling**
J. Woetzel ; R. Koch
1st European Conference on Visual Media Production (CVMP) 2004
Year: 2004
Pages: 245 - 254
Cited by: Papers (11) | Patents (2)
IET Conferences

Abstract (866 Kb)

Multi-camera real-time depth estimation with discontinuity handling on PC graphics hardware

J. Woetzel | R. Koch
Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.
Year: 2004 | Volume: 1

☐ **Multi-camera real-time depth estimation with discontinuity handling on PC graphics hardware**
J. Woetzel ; R. Koch
Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.
Year: 2004, Volume: 1
Pages: 741 - 744 Vol.1
Cited by: Papers (4)

IEEE Computer Graphics and Applications
Year: 2005 | Volume: 25 | Issue: 5

Hardware-assisted feature analysis and visualization of procedurally encoded multifield volumetric data
M. Weiler ; R. Botchen ; S. Stegmaier ; T. Ertl ; J. Huang ; Y. Jang ; D.S. Ebert ; K.P. Gaither
IEEE Computer Graphics and Applications
Year: 2005, Volume: 25, Issue: 5
Pages: 72 - 81
Cited by: Papers (14)
IEEE Journals & Magazines

Abstract (944 Kb)

Hardware-accelerated adaptive EWA volume splatting

Wei Chen | Liu Ren | M. Zwicker | H. Pfister
IEEE Visualization 2004
Year: 2004

Hardware-accelerated adaptive EWA volume splatting
Wei Chen; Liu Ren; M. Zwicker ; H. Pfister
IEEE Visualization 2004
Year: 2004
Pages: 67 - 74
Cited by: Papers (8)
IEEE Conferences

Abstract (1923 Kb)

Fast maximum intensity projections of large medical data sets by exploiting hierarchical memory architectures

G. Kiefer | H. Lehmann | J. Weese
IEEE Transactions on Information Technology in Biomedicine
Year: 2006 | Volume: 10 | Issue: 2

Fast maximum intensity projections of large medical data sets by exploiting hierarchical memory architectures
G. Kiefer ; H. Lehmann ; J. Weese
IEEE Transactions on Information Technology in Biomedicine
Year: 2006, Volume: 10, Issue: 2
Pages: 385 - 394
Cited by: Papers (6)
IEEE Journals & Magazines

Abstract (493 Kb)

Parallel hyperspectral image processing on commodity graphics hardware

J. Setoain | C. Tenllado | M. Prieto | D. Valencia | A. Plaza | J. Plaza
2006 International Conference on Parallel Processing Workshops (ICPPW'06)
Year: 2006

Parallel hyperspectral image processing on commodity graphics hardware
J. Setoain ; C. Tenllado ; M. Prieto ; D. Valencia ; A. Plaza ; J. Plaza
2006 International Conference on Parallel Processing Workshops (ICPPW'06)
Year: 2006
Pages: 8 pp. - 472
Cited by: Papers (4)
IEEE Conferences

Abstract (606 Kb)

« First < 1 > Next >>

IEEE Account	Purchase Details	Profile Information	Need Help?
» Change Username/Password	» Payment Options	» Communications Preferences	» US & Canada: +1 800 678 4333
» Update Address	» Order History	» Profession and Education	» Worldwide: +1 732 981 0060
	» View Purchased Documents	» Technical Interests	» Contact & Support

| About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

IEEE Account

Profile Information

Purchase Details

Need Help?

Other

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2018 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

US & Canada: +1 800 678 4333
Worldwide: +1 732 981 0060

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within results Show: Show All Results | Download PDFs | Per Page: Per Page 25 | Export | Alert Set Search Alerts | Search History

Displaying results 1-3 of 3 for ("GPU" AND ("simultaneous" NEAR/3 "transfer"))

Filters Applied: 1982 - 2006

☐ Journals & Magazines (2) ☐ Conferences (1)

Filter | Alert Set Search Alerts show: Show All Results | sort: Sort Relevance

Year

Single Year | Range

1982 2003

From 1982 To 2003

Author

Affiliation

Publication Title

Publisher

Conference Location

☐ **Bulk Power Area Reliability Evaluation Considering Propabilistic Transfer Capability**
Transmission Reliability Task Force
IEEE Transactions on Power Apparatus and Systems
Year: 1982 , Volume: PAS-101 , Issue: 9
Page s: 3551 - 3562
Cited by: Papers (7)
IEEE Journals & Magazines
Abstract (2094 Kb)

Bulk Power Area Reliability Evaluation Considering Propabilistic Transfer Capability

Transmission Reliability Task Force
IEEE Transactions on Power Apparatus and Systems
Year: 1982 | Volume: PAS-101 | Issue: 9

☐ **High-quality two-level volume rendering of segmented data sets on consumer graphics hardware**
M. Hadwiger ; C. Berger ; H. Hauser
IEEE Visualization, 2003. VIS 2003.
Year: 2003
Page s: 301 - 308
Cited by: Papers (42) | Patents (14)
IEEE Conferences
Abstract (451 Kb)

High-quality two-level volume rendering of segmented data sets on consumer graphics hardware

M. Hadwiger | C. Berger | H. Hauser
IEEE Visualization, 2003. VIS 2003.
Year: 2003

☐ **Top three technical problems-a survey report [of power system planning]**
P.P. Sener ; E.R. Greene
IEEE Transactions on Power Systems
Year: 1997 , Volume: 12 , Issue: 1
Page s: 230 - 244
Cited by: Papers (1)
IEEE Journals & Magazines
Abstract (1271 Kb)

Top three technical problems-a survey report [of power system planning]

P.P. Sener | E.R. Greene

| IEEE.org | IEEE *Xplore* Digital Library | IEEE-SA | IEEE Spectrum | More Sites | Cart (0) | Create Account | Personal Sign In

Access provided by:
United States Patent and Trademark Office
Sign Out

Browse My Settings Get Help

Search within results Show: Show All Results | Download PDFs | Per Page: Per Page 25 | Export | Alert Set Search Alerts | Search History

Displaying results 1-25 of 51 for ("GPU" AND ("concurrent" NEAR/3 "transfer"))

Conferences (42) Journals & Magazines (8) Early Access Articles (1)

Filter | Alert Set Search Alerts show: Show All Results | sort: Sort Relevance

Year

Single Year | Range

2008 2018

From 2008 To 2018

Author

Affiliation

Publication Title

Publisher

Supplemental Items

Conference Location

Performance improvement of CUDA applications by reducing CPU-GPU data transfer overhead
N. V. Sunitha ; K. Raju ; Niranjan N. Chiplunkar
2017 International Conference on Inventive Communication and Computational Technologies (ICICCT)
Year: 2017
Page s: 211 - 215
IEEE Conferences
Abstract (243 Kb)

Performance improvement of CUDA applications by reducing CPU-GPU data transfer overhead

N. V. Sunitha | K. Raju | Niranjan N. Chiplunkar
2017 International Conference on Inventive Communication and Computational Technologies (ICICCT)
Year: 2017

Moka: Model-based concurrent kernel analysis
Leiming Yu ; Xun Gong ; Yifan Sun ; Qianqian Fang ; Norm Rubin ; David Kaeli
2017 IEEE International Symposium on Workload Characterization (IISWC)
Year: 2017
Page s: 197 - 206
IEEE Conferences
Abstract (632 Kb)

Moka: Model-based concurrent kernel analysis

Leiming Yu | Xun Gong | Yifan Sun | Qianqian Fang | Norm Rubin | David Kaeli
2017 IEEE International Symposium on Workload Characterization (IISWC)
Year: 2017

GPU-Accelerated Multi-Profile Radiative Transfer Model for the Infrared Atmospheric Sounding Interferometer
Jarno Mielikainen ; Bormin Huang ; Hung-Lung Allen Huang
IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing
Year: 2011 , Volume: 4 , Issue: 3
Page s: 691 - 700
Cited by: Papers (32)
IEEE Journals & Magazines
Abstract (990 Kb)

GPU-Accelerated Multi-Profile Radiative Transfer Model for the Infrared Atmospheric Sounding Interferometer

Jarno Mielikainen | Bormin Huang | Hung-Lung Allen Huang
IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing

STIC
Scientific and Technical Information Center

EIC 2600

SEARCH REPORT

STIC Database Tracking Number: 571519

To: William Wood
Location: MDE-4C03
Art Unit: N/A
Monday, June 18, 2018

Case Serial Number: 15/808201

From: Kimberly Johnson
Location: EIC2600
KNX-8B59
Phone: (571)272-3493

kimberly.johnson@uspto.gov

Search Notes

US9189828B2

11/17/2015

Graphic processor based accelerator system and method
Granted

An accelerator system is implemented on an expansion card comprising a printed circuit board having (a) one or more graphics processing units (GPUs), (b) two or more associated memory banks (logically or physically partitioned), (c) a specialized controller, and (d) a local bus providing signal coupling compatible with the PCI industry standards. The controller handles most of the primitive operations to set up and control GPU computation. Thus, the computer's central processing unit (CPU) can be dedicated to other tasks. In this case a few controls (simulation start and stop signals from the CPU and the simulation completion signal back to CPU), GPU programs and input/output data are exchanged between CPU and the expansion card. Moreover, since on every time step of the simulation the results from the previous time step are used but not changed, the results are preferably transferred back to CPU in parallel with the computation.
Last viewed:
Bibliographic information

Publication US9189828B2 11/17/2015

Application US14147015 01/03/2014

Priority US14147015 01/03/2014

Current assignee NEURALA, INC.

Inventor Anatoli Gorchetchnikov

Legal representative Cooley LLP

Status Granted

Neurala, Inc. NEURALA
Current Assignees

Original	Standardized Normalized
NEURALA, INC.	NEURALA

Inventor(s)
Anatoli Gorchetchnikov
Heather Marie Ames
Massimiliano Versace

Fabrizio Santini

Legal representative(s)

Cooley LLP

Legal information

Status

Event	Date
Granted	11/17/2015
Filed	01/03/2014

Ownership

Date from	Date to	Name	Standardized	Normalized
02/21/2013		NEURALA, INC.	NEURALA	
10/12/2007	10/12/2007	GORCHETCHNIKOV, ANATOLI	GORCHETCHNIKOV ANATOLI	
10/12/2007	10/12/2007	AMES, HEATHER MARIE	AMES HEATHER MARIE	
10/12/2007	10/12/2007	VERSACE, MASSIMILIANO	VERSACE MASSIMILIANO	
10/12/2007	10/12/2007	SANTINI, FABRIZIO	SANTINI FABRIZIO	

Payment status

Date	Payment status	Description

Litigation

Date	Docket number	Comments	Judgement

Select Category
Cases 0

Results for: 9189828 *or* 9,189,828

Actions

Filters

Patent Law; Cases

Clear

Filter Results

Cases (0)

No documents found in **Patent Law; Cases**.

Try the following:

Remove one or more filters: **Patent Law; Cases**

Run your search again in Everything.

Modify your search (edit or remove some search terms or add synonyms).

Check for spelling errors.

Use the options available from the **Actions** menu above to modify your search.

Note: Some content may not be visible based on the restrictions of your subscription.

About LexisNexis® Privacy Policy Terms & Conditions Sign Out Copyright © 2018 LexisNexis. All rights reserved.

Select Category
Secondary... 0

Results for: 9189828 *or* 9,189,828

Actions

Filters

Patent Law; Law Reviews and Journals

Clear

Filter Results

Secondary Materials (0)

No documents found in **Patent Law; Law Reviews and Journals**.

Try the following:

Remove one or more filters: **Patent Law; Law Reviews and Journals**

Run your search again in Everything.

Modify your search (edit or remove some search terms or add synonyms).

Check for spelling errors.

Use the options available from the **Actions** menu above to modify your search.

Note: Some content may not be visible based on the restrictions of your subscription.

About LexisNexis® Privacy Policy Terms & Conditions Sign Out Copyright © 2018 LexisNexis. All rights reserved.

Select Category
News 9

Results for: 9189828 *or* 9,189,828

Actions

Filters

News

Clear

Filter Results

News (9)

Group Duplicates: Off

Sort by: Relevance

1. U.S. Patents Awarded to Inventors in Massachusetts (Nov. 20)

Targeted News Service | Nov 20, 2015 | 3014 words | Targeted News Service Targeted News Service

... The full-text of the patent can be found at http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=%2Fnetahtml%2FPTO%2Fsrchnum.htm&r=1&f=G&l=50&s1=**9,189,828**.PN.&OS=PN/**9,189,828**&RS=PN/**9,189,828**
Written by Kusum Sangma; edited by Jaya Anand. *** Johns ...
... Nov. 20 -- Neurala, Boston, has been assigned a patent (**9,189,828**) developed by four co-inventors for "graphic processor based accelerator system ...

2. Neurala Assigned Patent

Targeted News Service | Nov 20, 2015 | 151 words | Targeted News Service

... The full-text of the patent can be found at http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=%2Fnetahtml%2FPTO%2Fsrchnum.htm&r=1&f=G&l=50&s1=**9,189,828**.PN.&OS=PN/**9,189,828**&RS=PN/**9,189,828**
Written by Kusum Sangma; edited by Jaya Anand. For more ...
... Nov. 20 -- Neurala, Boston, has been assigned a patent (**9,189,828**) developed by four co-inventors for "graphic processor based accelerator system ...

3.

Head Line: US Patent granted to Neurala, Inc (Massachusetts) on November 17, 2015 titled as "Graphic processor based accelerator system and method"

Plus Patent News | Nov 17, 2015 | 252 words

... States Patent and Trademark Office has granted patent no. **9,189,828** on November 17, 2015, to Neurala, Inc (...

4. United States Patent for Graphic Processor Based Accelerator System and Method Issued to Neurala Inc

Global IP News. Information Technology Patent News | Nov 17, 2015 | 316 words

... Application Number: 14/147,015 Patent Publication Number: **9,189,828** International Patent Classification Codes: G06T 1/60 (20130101), G06F ...

5. United States Patent for Graphic Processor Based Accelerator System and Method Issued to Neurala Inc

Global IP News. Information Technology Patent News | Nov 17, 2015 | 314 words

... Technology Patent News Patent Application Number: 14/147,015 Patent Publication Number: **9,189,828** International Patent Classification Codes: G06T 1/60 (20130101), G06F 9/5027 (20130101), ...

6. 7 Relationship Milestones That Are Just As Meaningful As Marriage

Newstex Blogs | Aug 05, 2016 | 953 words | Brittany Wong

... http://www.huffingtonpost.com/natalie-brooke/getting-married-is-not-an-accomplishment_b_**9189828** .html?ncid=txtlnkusaolp00000603 [3]: ...

7. Ashley Csanady: Even on the podium in Rio, a woman's accomplishments are overshadowed by her marital status. Thanks, Hollywood

Postmedia Breaking News | Aug 15, 2016 | RIO 2016 | 1371 words

... and then babies, as the pinnacle of female accomplishment (http://www.huffingtonpost.com/natalie-brooke/getting-married-is-not-an-accomplishment_b_**9189828** .html), even in the face of Olympic glory. One ...

8. 10 Signs You're Not Ready To Get Married, According To Experts

Newstex Blogs | May 02, 2016 | 1189 words | Carolin Lehmann

... http://www.huffingtonpost.com/natalie-brooke/getting-married-is-not-an-accomplishment_b_**9189828** .html?ncid=txtlnkusaolp00000603 [2]: ...

☐ 9. ANOTHER ENERGY SAVING TIP FROM

GlobalAdSource (English) | Sep 26, 2007 | 64 words

... ANOTHER ENERGY SAVING TIP FROM ID **9189828** Description ORANGE SCREEN / Winterize / Male V/O- You don't ...

About LexisNexis® Privacy Policy Terms & Conditions Sign Out Copyright © 2018 LexisNexis. All rights reserved.

Lexis Advance® Research

More

Document: Head Line: US Patent granted to Neurala, Inc (Massachusetts) on November 17, 2015 title... Actions

FULL | KWIC™

Go to Search Document

3 of 9 | Results list

Head Line: US Patent granted to Neurala, Inc (Massachusetts) on November 17, 2015 titled as "Graphic processor based accelerator system and method"

Plus Patent News

November 17, 2015 Tuesday

Copyright 2015 Plus Media Solutions Private Limited All Rights Reserved

Length: 252 words

Dateline: Lahore

Body

ALEXANDRIA, Va: United States Patent and Trademark Office has granted patent no. **9,189,828** on November 17, 2015, to Neurala, Inc (Massachusetts) titled as "Graphic processor based accelerator system and method"

Inventors: Gorchetchnikov; Anatoli (Boston, MA); Ames; Heather Marie (South Boston, MA); Versace; Massimiliano (South Boston, MA); Santini; Fabrizio (Jamaica Plain, MA)

Assignee: Neurala, Inc. (Boston, MA) (Massachusetts)

According to the abstract released by the U.S. Patent & Trademark Office: "An accelerator system is implemented on an expansion card comprising a printed circuit board having (a) one or more graphics processing units (GPUs), (b) two or more associated memory banks (logically or physically partitioned), (c) a specialized controller, and (d) a local bus providing signal coupling compatible with the PCI industry standards. The controller handles most of the primitive operations to set up and control GPU computation. Thus, the computer's central processing unit (CPU) can be dedicated to other tasks. In this case a few controls (simulation start and stop signals from the CPU and the simulation completion signal back to CPU), GPU programs and input/output data are exchanged between CPU and the expansion card. Moreover, since on every time step of the simulation the results from the previous time step are used but not changed, the results are preferably transferred back to CPU in parallel with the computation."

The patent was filed on January 3, 2014 Application No. 14/147,015

Classification

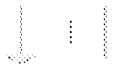
Language: ENGLISH

Subject: PATENTS (90%); COMPUTER GRAPHICS (77%)

Industry: COMPUTER GRAPHICS (77%); CIRCUIT BOARDS (70%)

Geographic: BOSTON, MA, USA (92%); MASSACHUSETTS, USA (95%); VIRGINIA, USA (79%); UNITED STATES (73%)

Load-Date: November 18, 2015



About LexisNexis® Privacy Policy Terms & Conditions Sign Out Copyright © 2018 LexisNexis All rights reserved.

LexisNexis® CourtLink®

Switch Client | My Briefcase | Order Runner Documents | Lexis Advance | Sign Out |
Welcome, Kim Johnson

Single Search - with Terms and Connectors

Enter keywords - Search multiple dockets & documents Search Info

My CourtLink | Search | Dockets & Documents | Track | Alert | Strategic Profiles | Breaking Complaints | My Account | Court Info

Search > Patent Search > Searching

Patent Search 8189828 6/ 19/ 2018

No cases found.

Return to Search

(Charges for search still apply)

LexisNexis® | About LexisNexis | Terms & Conditions | Pricing | Privacy | Customer Support - 1-888-311-1966 | Copyright © 2018 LexisNexis®. All rights reserved.



Search Strategy from ProQuest Dialog

June 19 2018 10:31

Search Strategy

Databases: LitAlert®

Set#	Searched for	Results
S1	pnum(us9189828)	0

NVIDIA Ex1002 Page 602 of 1081

Contact ProQuest
Copyright © 2018 ProQuest LLC. All rights reserved. - Terms and Conditions

PTO/SB/08a (07-09)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code2 (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
	1.	5,136,687	08-04-1992	Edelman et al.	
	2.	5,172,253	12-15-1992	Lynne	
	3.	5,388,206	02-07-1995	Poulton et al.	
	4.	6,336,051	01-01-2002	Pangels et al.	
	5.	6,647,508	11-11-2003	Zalewski et al.	
	6.	7,765,029	07-27-2010	Fleischer et al.	
	7.	7,861,060	12-01-2010	Nickolls et al.	
	8.	7,873,650	01-18-2011	Chapman et al.	
	9.	8,392,346	03-05-2013	Ueda et al.	
	10.	8,510,244	08-13-2013	Carson et al.	
	11.	8,648,867	02-11-2014	Gorchetchnikov et al.	
	12.	9,031,692	05-12-2015	Zhu	
	13.	9,177,246	11-03-2015	Buibas et al.	
	14.	9,626,566	04-18-2017	Versace et al.	
	15.	2002/0046271	04-18-2002	Huang	
	16.	2002/0050518	05-02-2002	Roustaei	
	17.	2003/0078754	04-24-2003	Hamza	
	18.	2004/0015334	01-22-2004	Ditlow et al.	
	19.	2005/0166042	07-28-2005	Evans	
	20.	2006/0184273	08-17-2006	Sawada et al.	
	21.	2007/0052713	03-08-2007	Chung et al.	
	22.	2007/0198222	08-23-2007	Schuster et al.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code2 (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
	23.	2007/0279429	12-06-2007	Ganzer	
	24.	2008/0258880	10-23-2008	Smith et al.	
	25.	2008/0033897	02-07-2008	Lloyd	
	26.	2008/0066065	03-13-2008	Kim et al.	
	27.	2009/0089030	04-02-2009	Sturrock et al.	
	28.	2010/0048242	02-25-2010	Rhoads et al.	
	29.	2011/0004341	01-06-2011	Sarvadevabhatla et al.	
	30.	2011/0173015	07-14-2011	Chapman et al.	
	31.	2012/0072215	03-22-2012	Yu et al.	
	32.	2012/0089552	04-12-2012	Chang et al.	
	33.	2012/0197596	08-02-2012	Comi	
	34.	2012/0316786	12-13-2012	Liu et al.	
	35.	2014/0052679	11-20-2014	Sinyavskiy et al.	
	36.	2014/0089232	03-27-2014	Buibas et al.	
	37.	2014/0019392	01-16-2014	Buibas et al.	
	38.	2014/0032461	01-30-2014	Weng	
	39.	2015/0127149	05-07-2015	Sinyavskiy et al.	
	40.	2015/0134232	05-14-2015	Robinson	
	41.	2016/0075017	03-17-2016	Laurent et al.	
	42.	2016/0082597	03-24-2016	Gorchetchnikov et al.	
	43.	2016/0096270	04-07-2016	Ibarz Gabardos et al.	
	44.	2016/0198000	07-07-2016	Gorchetchnikov et al.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

U.S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code2 (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
	45.	2017/0024877	01-26-2017	Versace et al.	
	46.	2017/0076194	03-16-2017	Versace et al.	
	47.	2017/0193298	07-06-2017	Versace et al.	

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Foreign Patent Document Country Code[3]-Number[4]-Kind Code[5] (*if known*)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages Or Relevant Figures Appear	T[6]
	48.	EP 1 224 622 B1	07-24-2002	Sun Microsystems, Inc.		
	49.	WO 2014/190208	11-27-2014	Neurala, Inc.		
	50.	WO 2014/204615	12-24-2014	Neurala, Inc.		
	51.	WO 2015/143173	09-24-2015	Neurala, Inc.		
	52.	WO 2016/014137	01-28-2016	Neurala, Inc.		

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	53.	Adelson, E. H., Anderson, C. H., Bergen, J. R., Burt, P. J., & Ogden, J. M. (1984). Pyramid methods in image processing. RCA engineer, 29(6), 33-41.	
/WHW/	54.	Aggarwal, Charu C, Hinneburg, Alexander, and Keim, Daniel A. On the surprising behavior of distance metrics in high dimensional space. Springer, 2001.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	55.	Ames, H, Versace, M., Gorchetchnikov, A., Chandler, B., Livitz, G., Léveillé, J., Mingolla, E., Carter, D., Abdalla, H., and Snider, G. (2012) Persuading computers to act more like brains. In Advances in Neuromorphic Memristor Science and Applications, Kozma, R.Pino,R., and Pazienza, G. (eds), Springer Verlag.	
/WHW/	56.	Ames, H. Mingolla, E., Sohail, A., Chandler, B., Gorchetchnikov, A., Léveillé, J., Livitz, G. and Versace, M. (2012) The Animat. IEEE Pulse, February 2012, 3(1), 47-50.	
/WHW/	57.	Artificial Intelligence As a Service. Invited talk, Defrag, Broomfield, CO, Nov 4-6 (2013).	
/WHW/	58.	Aryananda, L. (2006). Attending to learn and learning to attend for a social robot. Humanoids 06, pp. 618-623.	
/WHW/	59.	Baraldi, A. and Alpaydin, E. (1998). Simplified ART: A new class of ART algorithms. International Computer Science Institute, Berkeley, CA, TR-98-004, 1998.	
/WHW/	60.	Baraldi, A. and Alpaydin, E. (2002). Constructive feedforward ART clustering networks – Part I. IEEE Transactions on Neural Networks 13(3), 645–661.	
/WHW/	61.	Baraldi, A. and Parmiggiani, F. (1997). Fuzzy combination of Kohonen's and ART neural network models to detect statistical regularities in a random sequence of multi-valued input patterns. In International Conference on Neural Networks, IEEE.	
/WHW/	62.	Baraldi, Andrea and Alpaydin, Ethem. Constructive feedforward ART clustering networks—part II. IEEE Transactions on Neural Networks, 13(3):662–677, May 2002. ISSN 1045-9227. doi: 10.1109/tnn.2002.1000131. URL http://dx.doi.org/10.1109/tnn.2002.1000131.	
/WHW/	63.	Bengio, Y., Courville, A., & Vincent, P. Representation learning: A review and new perspectives, IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 35 Issue 8, August 2013, Pages 1798-1828 .	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT (Use as many sheets as necessary)	Complete if Known	
	Application Number	15/808,201
	Filing Date	11-09-2017
	First Named Inventor	Anatoli Gorchetchnikov
	Art Unit	2611
	Examiner Name	Not Yet Assigned
	Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	64.	Berenson, D. et al., A robot path planning framework that learns from experience, 2012 International Conference on Robotics and Automation, 2012, 9 pages [retrieved from the internet] URL:http://users.wpi.edu/-dberenson/lightning.pdf.	
/WHW/	65.	Bernhard, F., and Keriven, R. (2005). Spiking Neurons on GPUs. Tech. Rep. 05-15, Ecole Nationale des Ponts et Chauss´es, 8 pages.	
/WHW/	66.	Besl, P. J., & Jain, R. C. (1985). Three-dimensional object recognition. ACM Computing Surveys (CSUR), 17(1), 75-145.	
/WHW/	67.	Bohn, C.-A. Kohonen. (1998). Feature Mapping Through Graphics Hardware. In Proceedings of 3rd Int. Conference on Computational Intelligence and Neurosciences, 4 pages.	
/WHW/	68.	Bradski, G., & Grossberg, S. (1995). Fast-learning VIEWNET architectures for recognizing three-dimensional objects from multiple two-dimensional views. Neural Networks, 8 (7-8), 1053-1080.	
	~~69.~~	~~Brain-inspired computing. Invited keynote address, Bionetics 2010, Boston, MA, USA.~~	
/WHW/	70.	Canny, J.A. (1986). Computational Approach To Edge Detection, IEEE Trans. Pattern Analysis and Machine Intelligence, 8(6):679–698.	
/WHW/	71.	Carpenter, G.A. and Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. Computer Vision, Graphics, and Image Processing 37, 54–115.	
/WHW/	72.	Carpenter, G.A., and Grossberg, S. (1995). Adaptive resonance theory (ART). In M. Arbib (Ed.), The handbook of brain theory and neural networks. (pp. 79–82). Cambridge, M.A.: MIT press.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	73.	Carpenter, G.A., Grossberg, S. and Rosen, D.B. (1991). Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system. Neural Networks 4, 759–771.	
/WHW/	74.	Carpenter, Gail A and Grossberg, Stephen. The art of adaptive pattern recognition by a self-organizing neural network. Computer, 21(3):77–88, 1988.	
/WHW/	75.	Coifman, R.R. and Maggioni, M. Diffusion wavelets. Applied and Computational Harmonic Analysis, 21(1):53–94, 2006.	
/WHW/	76.	Coifman, R.R., Lafon, S., Lee, A.B., Maggioni, M., Nadler, B., Warner, F., and Zucker, S.W. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. Proceedings of the National Academy of Sciences of the United States of America, 102(21):7426, 2005.	
/WHW/	77.	Davis, C. E. 2005. Graphic Processing Unit Computation of Neural Networks. Master's thesis, University of New Mexico, Albuquerque, NM, 121 pages.	
/WHW/	78.	Dosher, B.A., and Lu, Z.L. (2010). Mechanisms of perceptual attention in precuing of location. Vision Res., 40(10–12). 1269–1292.	
/WHW/	79.	Ellias, S. A., and GROSSBERG, S. (1975). Pattern formation, contrast control and oscillations in the short term memory of shunting on-center off-surround networks. Biol Cybern 20, pp. 69–98.	
/WHW/	80.	Extended European Search Report and Written Opinion mailed June 1, 2017 from European Application No. 14813864.7, 10 pages (Atty. Docket No. NRLA-002/01EP).	
/WHW/	81.	Extended European Search Report and Written Opinion mailed October 12, 2017 from European Application No. 14800348.6, 12 pages (Atty. Docket No. NRLA-003/01EP).	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	82.	Extended European Search Report and Written Opinion October 23, 2017 from European Application No. 15765396.5, 8 pages (Atty. Docket No. NRLA-004/01EP).	
/WHW/	83.	Fazl, A., Grossberg, S., and Mingolla, E. (2009). View-invariant object category learning, recognition, and search: How spatial and object attention are coordinated using surface-based attentional shrouds. Cognitive Psychology 58, 1–48.	
/WHW/	84.	Földiák, P. (1990). Forming sparse representations by local anti-Hebbian learning, Biological Cybernetics, vol. 64, pp. 165-170.	
/WHW/	85.	Friston K., Adams R., Perrinet L., & Breakspear M. (2012). Perceptions as hypotheses: saccades as experiments. Frontiers in Psychology, 3 (151), 1-20.	
/WHW/	86.	Galbraith, B.V, Guenther, F.H., and Versace, M. (2015) A neural network-based exploratory learning and motor planning system for co-robots.Frontiers in Neuroscience, in press.	
/WHW/	87.	George, D. and Hawkins, J. (2009). Towards a mathematical theory of cortical micro-circuits. PLoS Computational Biology 5(10), 1–26.	
/WHW/	88.	Georgeii, J., and Westermann, R. (2005). Mass-spring systems on the GPU. Simulation Modelling Practice and Theory 13, pp. 693–702.	
/WHW/	89.	Gorchetchnikov A., Hasselmo M.E. (2005). A biophysical implementation of a bidirectional graph search algorithm to solve multiple goal navigation tasks. Connection Science, 17(1-2), pp. 145-166.	
/WHW/	90.	Gorchetchnikov A., Hasselmo M.E. (2005). A simple rule for spike-timing-dependent plasticity: local influence of AHP current. Neurocomputing, 65-66, pp. 885-890.	
/WHW/	91.	Gorchetchnikov A., Versace, M., Hasselmo M.E. (2005). A Model of STDP Based on Spatially and Temporally Local Information: Derivation and Combination with Gated Decay. Neural Networks, 18, pp. 458-466.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	92.	Gorchetchnikov A., Versace, M., Hasselmo M.E. (2005). Spatially and temporally local spike-timing-dependent plasticity rule. In: Proceedings of the International Joint Conference on Neural Networks, number 1568 in IEEE CD-ROM Catalog Number: 05CH37662C, pp. 390-396.	
/WHW/	93.	Gorchetchnikov A. An Approach to a Biologically Realistic Simulation of Natural Memory. Master's thesis, Middle Tennessee State University, Murfreesboro, TN, 70 pages.	
/WHW/	94.	Grossberg, S. (1973). Contour enhancement, short-term memory, and constancies in reverberating neural networks. Studies in Applied Mathematics 52, 213–257.	
/WHW/	95.	Grossberg, S., and Huang, T.R. (2009). ARTSCENE: A neural system for natural scene classification. Journal of Vision, 9 (4), 6.1–19. doi:10.1167/9.4.6	
/WHW/	96.	Grossberg, S., and Versace, M. (2008) Spikes, synchrony, and attentive learning by laminar thalamocortical circuits. Brain Research, 1218C, 278-312 [Authors listed alphabetically].	
/WHW/	97.	Hagen, T.R., Hjelmervik, J., Lie, K.-A., Natvig, J. and Ofstad Henriksen, M. (2005). Visual simulation of shallow-water waves. Simulation Modelling Practice and Theory 13, pp. 716–726.	
/WHW/	98.	Hasselt, Hado Van. Double q-learning. In Advances in Neural Information Processing Systems, pp. 2613–2621, 2010.	
/WHW/	99.	Hinton, G. E., Osindero, S., and Teh, Y. (2006). A fast learning algorithm for deep belief nets. Neural Computation, 18, 1527–1554.	
/WHW/	100.	Hodgkin, A.L., and Huxley, A.F. (1952). Quantitative description of membrane current and its application to conduction and excitation in nerve. J Physiol 117, pp. 500–544.	
/WHW/	101.	Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. In Proc Natl Acad Sci USA, vol. 79, pp. 2554–2558.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	102.	Ilie, A. (2002). Optical character recognition on graphics hardware. Tech. Rep. integrative paper, UNCCH, Department of Computer Science, 9 pages.	
/WHW/	103.	International Preliminary Report on Patentability in related PCT Application No. PCT/US2014/039162 filed 22 May 2014, mailed 24 November 2015, 7 pages.	
/WHW/	104.	International Preliminary Report on Patentability in related PCT Application No. PCT/US2014/039239 filed 22 May 2014, mailed 24 November 2015, 8 pages.	
/WHW/	105.	International Preliminary Report on Patentability mailed November 8, 2016 from International Application No. PCT/US2015/029438, 7 pages (Atty. Docket No. NRLA-006/01WO).	
/WHW/	106.	International Search Report and Written Opinion mailed February 18, 2015 from International Application No. PCT/US2014/039162, 12 pages (Atty. Docket No. NRLA-002/01WO).	
/WHW/	107.	International Search Report and Written Opinion mailed February 23, 2016 from International Application No. PCT/US2015/029438, 11 pages (Atty. Docket No. NRLA-006/01WO).	
/WHW/	108.	International Search Report and Written Opinion mailed July 6, 2017 from International Application No. PCT/US2017/029866, 12 pages (Atty. Docket No. NRLA-009/01WO).	
/WHW/	109.	International Search Report and Written Opinion mailed November 26, 2014 from International Application No. PCT/US2014/039239, 14 pages (Atty. Docket No. NRLA-003/01WO).	
/WHW/	110.	International Search Report and Written Opinion mailed September 15, 2015 from International Application No. PCT/US2015/021492, 9 pages (Atty. Docket No. NRLA-004/01WO).	
/WHW/	111.	Itti, L., and Koch, C. (2001). Computational modelling of visual attention. Nature Reviews Neuroscience, 2 (3), 194-203.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	112.	Itti, L., Koch, C., and Niebur, E. (1998). A Model of Saliency-Based Visual Attention for Rapid Scene Analysis, 1–6.	
/WHW/	113.	Jarrett, K., Kavukcuoglu, K., Ranzato, M. A., & LeCun, Y. (2009, September). What is the best multi-stage architecture for object recognition?. In Computer Vision, 2009 IEEE 12th International Conference on (pp. 2146-2153). IEEE.	
/WHW/	114.	Kipfer, P., Segal, M., and Westermann, R. (2004). UberFlow: A GPU-Based Particle Engine. In Proceedings of the SIGGRAPH/Eurographics Workshop on Graphics Hardware 2004, pp. 115–122.	
/WHW/	115.	Kolb, A., L. Latta, and C. Rezk-Salama. (2004). "Hardware-Based Simulation and Collision Detection for Large Particle Systems." In Proceedings of the SIGGRAPH/Eurographics Workshop on Graphics Hardware 2004, pp. 123–131.	
/WHW/	116.	Kompella, Varun Raj, Luciw, Matthew, and Schmidhuber, J¨urgen. Incremental slow feature analysis: Adaptive low-complexity slow feature updating from high-dimensional input streams. Neural Computation, 24(11):2994–3024, 2012.	
/WHW/	117.	Kowler, E. (2011). Eye movements: The past 25years. Vision Research, 51(13), 1457–1483. doi:10.1016/j.visres.2010.12.014	
/WHW/	118.	Larochelle H., & Hinton G. (2012). Learning to combine foveal glimpses with a third-order Boltzmann machine. NIPS 2010, 1243-1251.	
/WHW/	119.	LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010, May). Convolutional networks and applications in vision. In Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on (pp. 253-256). IEEE.	
/WHW/	120.	Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. Nature, 401(6755):788-791.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT (Use as many sheets as necessary)	Complete if Known	
	Application Number	15/808,201
	Filing Date	11-09-2017
	First Named Inventor	Anatoli Gorchetchnikov
	Art Unit	2611
	Examiner Name	Not Yet Assigned
	Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	121.	Lee, D. D., and Seung, H. S. (1997). "Unsupervised learning by convex and conic coding." Advances in Neural Information Processing Systems, 9.	
/WHW/	122.	Legenstein, R., Wilbert, N., and Wiskott, L. Reinforcement learning on slow features of high-dimensional input streams. PLoS Computational Biology, 6(8), 2010. ISSN 1553-734X.	
/WHW/	123.	Léveillé, J., Ames, H., Chandler, B., Gorchetchnikov, A., Mingolla, E., Patrick, S., and Versace, M. (2010) Learning in a distributed software architecture for large-scale neural modeling. BIONETICS10, Boston, MA, USA.	
	124.	Livitz G., Versace M., Gorchetchnikov A., Vasilkoski Z., Ames H., Chandler B., Leveille J. and Mingolla E. (2011) Scalable adaptive brain-like systems. The Neuromorphic Engineer, :	
		10.2417/1201101.003500 February 2011.	
/WHW/	125.	Livitz, G., Versace, M., Gorchetchnikov, A., Vasilkoski, Z., Ames, H., Chandler, B., Léveillé, J., Mingolla, E., Snider, G., Amerson, R., Carter, D., Abdalla, H., and Qureshi, S. (2011) Visually-Guided Adaptive Robot (ViGuAR). Proceedings of the International Joint Conference on Neural Networks (IJCNN) 2011, San Jose, CA, USA.	
/WHW/	126.	Lowe, D.G.(2004). Distinctive Image Features from Scale-Invariant Keypoints. Journal International Journal of Computer Vision archive Volume 60, 2, 91 – 110.	
/WHW/	127.	Lu, Z.L., Liu, J., and Dosher, B.A. (2010) Modeling mechanisms of perceptual learning with augmented Hebbian re- weighting. Vision Research, 50(4). 375–390.	
/WHW/	128.	Mahadevan, S. Proto-value functions: Developmental reinforcement learning. In Proceedings of the 22nd international conference on Machine learning, pp. 553–560. ACM, 2005.	
/WHW/	129.	Meuth, J.R. and Wunsch, D.C. (2007) A SURVEY OF NEURAL COMPUTATION ON GRAPHICS PROCESSING HARDWARE. 22nd IEEE International Symposium on Intelligent Control, Part of IEEE Multi-conference on Systems and Control, Singapore, 1-3 October 2007, 5 pages.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	130.	Mishkin M, Ungerleider LG. (1982). "Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys," Behav Brain Res, 6 (1): 57–77.	
/WHW/	131.	Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Rusu, Andrei A, Veness, Joel, Bellemare, Marc G, Graves, Alex, Riedmiller, Martin, Fidjeland, Andreas K, Ostrovski, Georg, et al. Human-level control through deep reinforcement learning. Nature, 518(7540):529–533, 25 February 2015.	
/WHW/	132.	Moore, Andrew W and Atkeson, Christopher G. Prioritized sweeping: Reinforcement learning with less data and less time. Machine Learning, 13(1):103–130, 1993.	
/WHW/	133.	Najemnik, J., and Geisler, W. (2009). Simple summation rule for optimal fixation selection in visual search. Vision Research. 49, 1286–1294.	
/WHW/	134.	Notice of Allowance mailed 07/27/2016 from U.S. Application Serial No. 14/662,657 (Atty. Docket No. NRLA-005/01US).	
/WHW/	135.	Notice of Allowance mailed 12/16/2016 from U.S. Application Serial No. 14/662,657 (Atty. Docket No. NRLA-005/01US).	
/WHW/	136.	Oh, K.-S., and Jung, K. (2004). GPU implementation of neural networks. Pattern Recognition 37, pp. 1311–1314.	
/WHW/	137.	Oja, E. (1982). Simplified neuron model as a principal component analyzer. Journal of Mathematical Biology 15(3), 267–273.	
/WHW/	138.	Partial Supplementary European Search Report dated July 4, 2017 from European Application No. 14800348.6, 13 pages (Atty. Docket No. NRLA-003/01EP).	
/WHW/	139.	Raijmakers, M.E.J., and Molenaar, P. (1997). Exact ART: A complete implementation of an ART network Neural networks 10 (4), 649-669.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	140.	Ranzato, M. A., Huang, F. J., Boureau, Y. L., & Lecun, Y. (2007, June). Unsupervised learning of invariant feature hierarchies with applications to object recognition. In Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on (pp. 1-8). IEEE.	
/WHW/	141.	Raudies, F., Eldridge, S., Joshi, A., and Versace, M. (20 August 2014). Learning to navigate in a virtual world using optic flow and stereo disparity signals. Artificial Life and Robotics, DOI 10.1007/s10015-014-0153-1.	
/WHW/	142.	Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. Nature Neuroscience, 2 (11), 1019-1025.	
/WHW/	143.	Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. Nature neuroscience, 3, 1199-1204.	
/WHW/	144.	Rolfes, T. (2004). Artificial Neural Networks on Programmable Graphics Hardware. In Game Programming Gems 4, A. Kirmse, Ed. Charles River Media, Hingham, MA, pp. 373–378.	
/WHW/	145.	Rublee, E., Rabaud, V., Konolige, K., & Bradski, G. (2011). ORB: An efficient alternative to SIFT or SURF. In IEEE International Conference on Computer Vision (ICCV) 2011, 2564–2571.	
/WHW/	146.	Ruesch, J. et al. (2008). Multimodal Saliency-Based Bottom-Up Attention A Framework for the Humanoid Robot iCub. 2008 IEEE International Conference on Robotics and Automation, pp. 962-965.	
/WHW/	147.	Rumelhart D., Hinton G., and Williams, R. (1986). Learning internal representations by error propagation. In Parallel distributed processing: explorations in the microstructure of cognition, vol. 1, MIT Press.	
/WHW/	148.	Rumpf, M. and Strzodka, R. Graphics processor units: New prospects for parallel computing. In Are Magnus Bruaset and Aslak Tveito, editors, Numerical Solution of Partial Differential Equations on Parallel Computers, volume 51 of Lecture Notes in Computational Science and Engineering, pp. 89–134. Springer, 2005.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT
(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
	149.	Salakhutdinov, R., & Hinton, G. E. (2009). Deep boltzmann machines. In International Conference on Artificial Intelligence and Statistics (pp. 448-455).	
/WHW/	150.	Schaul, Tom, Quan, John, Antonoglou, Ioannis, and Silver, David. Prioritized experience replay. arXiv preprint arXiv: 1511.05952, 18 November 2015.	
/WHW/	151.	Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990–2010). Autonomous Mental Development, IEEE Transactions on, 2(3), 230-247.	
/WHW/	152.	Schmidhuber, Jürgen. Curious model-building control systems. In Neural Networks, 1991. 1991 IEEE International Joint Conference on, pp. 1458–1463. IEEE, 1991.	
/WHW/	153.	Seibert, M., & Waxman, A.M. (1992). Adaptive 3-D Object Recognition from Multiple Views. IEEE Transactions on Pattern Analysis and Machine Intelligence, 14 (2), 107-124.	
/WHW/	154.	Sherbakov, L. and Versace, M. (2014) Computational principles for an autonomous active vision system. Ph.D., BOSTON UNIVERSITY, http://search.proquest.com/docview/1558856407	
	155.	Sherbakov, L. et al. (2012). CogEye: from active vision to context identification, youtube, retrieved from the Internet on 10-10-2017: URL://www.youtube.com/watch?v=i5PQk962B1k, 1 page.	
	156.	Sherbakov, L. et al. (2013). CogEye: system diagram module brain area function algorithm approx # neurons, retrieved from the Internet on 10-12-2017: URL://http://www-labsticc.univ-ubs.fr/~coussy/neucomp2013/index_fichiers/material/posters/NeuComp2013_final56x36.pdf, 1 page.	
/WHW/	157.	Sherbakov, L., Livitz, G., Sohail, A., Gorchetchnikov, A., Mingolla, E., Ames, H., and Versace, M. (2013a) CogEye: An online active vision system that disambiguates and recognizes objects. NeuComp 2013.	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.
[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT
(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	158.	Sherbakov, L., Livitz, G., Sohail, A., Gorchetchnikov, A., Mingolla, E., Ames, H., and Versace, M (2013b) A computational model of the role of eye-movements in object disambiguation. Cosyne, Feb 28 - Mar 3, 2013. Salt Lake City, UT, USA.	
/WHW/	159.	Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. In D. E.	
	160.	Snider, G., Amerson, R., Carter, D., Abdalla, H., Qureshi, S., Laveille, J., Versace, M., Ames, H., Patrick, S., Chandler, B., Gorchetchnikov, A., and Mingolla, E. (2010) Adaptive Computation with Memristive Memory. IEEE Computer, Volume 44(2).	
/WHW/	161.	Spratling, M. W. (2008). Predictive coding as a model of biased competition in visual attention. Vision Research, 48(12):1391-1408.	
/WHW/	162.	Spratling, M. W. (2012). Unsupervised learning of generative and discriminative weights encoding elementary image components in a predictive coding model of cortical function. Neural Computation, 24(1):60-103.	
/WHW/	163.	Spratling, M. W., De Meyer, K., and Kompass, R. (2009). Unsupervised learning of overlapping image components using divisive input modulation. Computational intelligence and neuroscience.	
/WHW/	164.	Sprekeler, H. On the relation of slow feature analysis and laplacian eigenmaps. Neural Computation, pp. 1–16, 2011.	
/WHW/	165.	Sutton, Richard S and Barto, Andrew G. Reinforcement learning: An introduction. MIT press, 1998.	
	~~166.~~	~~TEDx Fulbright, invited talk, Washington DC, 4/5/2014.~~	

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.
[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT LIST BY APPLICANT

(Use as many sheets as necessary)

Complete if Known	
Application Number	15/808,201
Filing Date	11-09-2017
First Named Inventor	Anatoli Gorchetchnikov
Art Unit	2611
Examiner Name	Not Yet Assigned
Attorney Docket Number	NRLA-001/03US 320730-2040

NON PATENT LITERATURE DOCUMENTS

Examiner's Initials	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title ofthe item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
/WHW/	167.	Tong, F., Ze-Nian Li, (1995). Reciprocal-wedge transform for space-variant sensing," Pattern Analysis and Machine Intelligence, IEEE Transactions on , vol.17, no.5, pp.500 - 51. doi: 10.1109/34.391393	
/WHW/	168.	Torralba, A., Oliva, A., Castelhano, M.S., Henderson, J.M. (2006). Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. Psychological Review, 113(4).766–786.	
/WHW/	169.	Van Hasselt, Hado, Guez, Arthur, and Silver, David. Deep reinforcement learning with double q-learning. arXiv preprint arXiv: 1509.06461, 22 September 2015.	
/WHW/	170.	Versace, M. (2006) From spikes to interareal synchrony: how attentive matching and resonance control learning and information processing by laminar thalamocortical circuits. NSF Science of Learning Centers PI Meeting, Washington, DC, USA.	
/WHW/	171.	Versace, M., (2010) Open-source software for computational neuroscience: Bridging the gap between models and behavior. In Horizons in Computer Science Research,Volume 3.	
/WHW/	172.	Versace, M., Ames, H., Léveillé, J., Fortenberry, B., and Gorchetchnikov, A. (2008) KInNeSS: A modular framework for computational neuroscience. Neuroinformatics, 2008 Winter; 6(4):291-309. Epub 2008 Aug 10.	
/WHW/	173.	Versace, M., and Chandler, B. (2010) MoNETA: A Mind Made from Memristors. IEEE Spectrum, December 2010.	
/WHW/	174.	Webster, Bachevalier, Ungerleider (1994). Connections of IT areas TEO and TE with parietal and frontal cortex in macaque monkeys. Cerebal Cortex, 4(5), 470-483.	
/WHW/	175.	Wiskott, Laurenz and Sejnowski, Terrence. Slow feature analysis: Unsupervised learning of invariances. Neural Computation, 14(4):715–770, 2002.	

156249486 v1

Examiner Signature:	/William H. Wood/	Date Considered	09/18/2018

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached. This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Anatoli Gorchechnikov

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT/RESPONSE TO OFFICE ACTION

In response to the Office Action dated November 20, 2018, to which the deadline for responding is March 20, 2019, with a one-month extension of time, Applicant submits the following Amendments and Remarks, and respectfully requests reconsideration of the application in view thereof.

Amendments to the Claims are reflected in the listing of the claims which begins on page 2 of this paper.

Remarks begin on page 23 of this paper.

Doc code: IDS
Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (02-18)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	10083523		2018-09-25	Versace et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	15808201
	Filing Date	2017-11-09
	First Named Inventor	Anatoli Gorchechnikov
	Art Unit	3992
	Examiner Name	William H WOOD
	Attorney Docket Number	NRLA-001US03

	1	Brain-inspired computing. Invited keynote address, Bionetics 2010, Boston, MA, USA. 1 page.
	2	Livitz G., Versace M., Gorchetchnikov A., Vasilkoski Z., Ames H., Chandler B., Leveille J. and Mingolla E. (2011) Adaptive, brain-like systems give robots complex behaviors, The Neuromorphic Engineer, : 10.2417/1201101.003500 February 2011. 3 pages.
	3	Salakhutdinov, R., & Hinton, G. E. (2009). Deep boltzmann machines. In International Conference on Artificial Intelligence and Statistics (pp. 448-455).
	4	SHERBAKOV, L. et al. 2012. CogEye: from active vision to context identification, youtube, retrieved from the Internet on 10-10-2017: URL://www.youtube.com/watch?v=i5PQk962B1k, 1 page.
	5	SHERBAKOV, L. et al. 2013. CogEye: system diagram module brain area function algorithm approx # neurons, retrieved from the Internet on 10-12-2017: URL://http://www-labsticc.univ-ubs.fr/~coussy/neucomp2013/index_fichiers/material/posters/NeuComp2013_final56x36.pdf, 1 page.
	6	Snider, Greg, et al. "From synapses to circuitry: Using memristive memory to explore the electronic brain." IEEE Computer, Volume 44(2). (2011): 21-28.
	7	VERSACE, TEDx Fulbright, Invited talk, Washington DC, 4/5/2014. 30 pages.

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Christopher Max Colice/	Date (YYYY-MM-DD)	2019-03-20
Name/Print	Christopher Max Colice	Registration Number	65634

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Doc Code: REIS.DECL
Document Description: Reissue Declaration Filed In Accordance With MPEP 1414

PTO/AIA/05 (06-12)
Approved for use through 08/31/2013. OMB 0651-0033
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

REISSUE APPLICATION DECLARATION BY THE INVENTOR	Docket Number (Optional) NRLA-001US03

I hereby declare that:
Each inventor's residence and mailing address are stated below next to their name.
I believe I am the original inventor or an original joint inventor of the subject matter which is described and claimed in patent number 9,189,828, granted November 17, 2015 and for which a reissue patent is sought on the invention titled Graphic Processor Based Accelerator System and Method

the specification of which

☐ is attached hereto.

☑ was filed on 11/9/2017 as reissue application number 15/808,201.

The above-identified application was made or authorized to be made by me.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

I believe the original patent to be wholly or partly inoperative or invalid, for the reasons described below. (Check all boxes that apply.)

☐ by reason of a defective specification or drawing.

☑ by reason of the patentee claiming more or less than he had the right to claim in the patent.

☑ by reason of other errors.

At least one error upon which reissue is based is described below. If the reissue is a broadening reissue, a claim that the application seeks to broaden must be identified:

Reissue is requested because of antecedent basis errors in issued claim 1 of US 9,189,828. Specifically, claim 1 lacks antecedent basis for "the plurality of computations," which is being amended to recite "the sequence of computations" to correct this error. Reissue is also requested because Applicant erred by claiming less than it had a right to claim in US 9,189,828. Claim 21, for example, corrects this error because it does not recite "a central processing unit" as in claim 1. And claim 79, for example, corrects Applicant's error by claiming the "user interaction stream" that appears in the embodiment of FIG. 3 of US 9,189,828 and has not been claimed to date.

This collection of information is required by 37 CFR 1.175. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 36 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/05 (06-12)
Approved for use through 08/31/2013. OMB 0651-0033
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

(REISSUE APPLICATION DECLARATION BY THE INVENTOR, page 2)

Docket Number (Optional): NRLA-001US03

Note: To appoint a power of attorney, use form PTO/AIA/81.

Correspondence Address: Direct all communications about the application to:

☑ The address associated with Customer Number: 149890

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

Legal name of sole or first inventor (*E.g.*, Given Name (first and middle (if any) and Family Name or Surname)

Heather Marie Ames

Inventor's Signature	Date (Optional)

Residence: City	State	Country
Milton	MA	US

Mailing Address

13 Myers Lane

City	State	Zip	Country
Milton	MA	02186	US

☐ Additional joint inventors are named on the _____ supplemental sheet(s) PTO/AIA/10 attached hereto.

NVIDIA Ex1002 Page 626 of 1081

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Doc Code: REIS.DECL
Document Description: Reissue Declaration Filed In Accordance With MPEP 1414

PTO/AIA/05 (08-12)
Approved for use through 08/31/2013. OMB 0651-0033
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

REISSUE APPLICATION DECLARATION BY THE INVENTOR	Docket Number (Optional) NRLA-001US03

I hereby declare that:
Each inventor's residence and mailing address are stated below next to their name.
I believe I am the original inventor or an original joint inventor of the subject matter which is described and claimed in patent number 9,189,828 granted November 17, 2015 and for which a reissue patent is sought on the invention titled Graphic Processor Based Accelerator System and Method

the specification of which

☐ is attached hereto.

☑ was filed on 11/9/2017 as reissue application number 15/808,201.

The above-identified application was made or authorized to be made by me.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

I believe the original patent to be wholly or partly inoperative or invalid, for the reasons described below. (Check all boxes that apply.)

☐ by reason of a defective specification or drawing.

☑ by reason of the patentee claiming more or less than he had the right to claim in the patent.

☑ by reason of other errors.

At least one error upon which reissue is based is described below. If the reissue is a broadening reissue, a claim that the application seeks to broaden must be identified:

Reissue is requested because of antecedent basis errors in issued claim 1 of US 9,189,828. Specifically, claim 1 lacks antecedent basis for "the plurality of computations," which is being amended to recite "the sequence of computations" to correct this error. Reissue is also requested because Applicant erred by claiming less than it had a right to claim in US 9,189,828. Claim 21, for example, corrects this error because it does not recite "a central processing unit" as in claim 1. And claim 79, for example, corrects Applicant's error by claiming the "user interaction stream" that appears in the embodiment of FIG. 3 of US 9,189,828 and has not been claimed to date.

This collection of information is required by 37 CFR 1.175. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/05 (06-12)
Approved for use through 08/31/2013. OMB 0651-0033
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

(REISSUE APPLICATION DECLARATION BY THE INVENTOR, page 2)

Docket Number (Optional): NRLA-001US03

Note: To appoint a power of attorney, use form PTO/AIA/81.

Correspondence Address: Direct all communications about the application to:

☑ The address associated with Customer Number: 149890

OR

☐ Firm or individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

Legal name of sole or first inventor (*E.g.*, Given Name (first and middle (if any) and Family Name or Surname)

Anatoli Gorchetchnikov

Inventor's Signature		Date (Optional)	
Residence: City	State	Country	
Newton	MA	US	
Mailing Address			
15 Jefferson St.			
City	State	Zip	Country
Newton	MA	02458	US

☐ Additional joint inventors are named on the _____ supplemental sheet(s) PTO/AIA/10 attached hereto.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Doc Code: REIS.DECL
Document Description: Reissue Declaration Filed In Accordance With MPEP 1414

PTO/AIA/05 (06-12)
Approved for use through 08/31/2013. OMB 0651-0033
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

REISSUE APPLICATION DECLARATION BY THE INVENTOR	Docket Number (Optional) NRLA-001US03

I hereby declare that:
Each inventor's residence and mailing address are stated below next to their name.
I believe I am the original inventor or an original joint inventor of the subject matter which is described and claimed in patent number 9,189,828, granted November 17, 2015 and for which a reissue patent is sought on the invention titled Graphic Processor Based Accelerator System and Method,

the specification of which

☐ is attached hereto.

☑ was filed on 11/9/2017 as reissue application number 15/808,201.

The above-identified application was made or authorized to be made by me.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

I believe the original patent to be wholly or partly inoperative or invalid, for the reasons described below. (Check all boxes that apply.)

☐ by reason of a defective specification or drawing.

☑ by reason of the patentee claiming more or less than he had the right to claim in the patent.

☑ by reason of other errors.

At least one error upon which reissue is based is described below. If the reissue is a broadening reissue, a claim that the application seeks to broaden must be identified:

Reissue is requested because of antecedent basis errors in issued claim 1 of US 9,189,828. Specifically, claim 1 lacks antecedent basis for "the plurality of computations," which is being amended to recite "the sequence of computations" to correct this error. Reissue is also requested because Applicant erred by claiming less than it had a right to claim in US 9,189,828. Claim 21, for example, corrects this error because it does not recite "a central processing unit" as in claim 1. And claim 79, for example, corrects Applicant's error by claiming the "user interaction stream" that appears in the embodiment of FIG. 3 of US 9,189,828 and has not been claimed to date.

This collection of information is required by 37 CFR 1.175. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/05 (06-12)
Approved for use through 08/31/2013. OMB 0651-0033
U.S. Patent and Trademark Office, U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

(REISSUE APPLICATION DECLARATION BY THE INVENTOR, page 2)

Docket Number (Optional): NRLA-001US03

Note: To appoint a power of attorney, use form PTO/AIA/81.

Correspondence Address: Direct all communications about the application to:

☑ The address associated with Customer Number: 149890

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

Legal name of sole or first inventor (*E.g.*, Given Name (first and middle (if any) and Family Name or Surname)

Massimiliano Versace

Inventor's Signature		Date (Optional): 2/13/13
Residence: City: Milton	State: MA	Country: US

Mailing Address

13 Myers Lane

City	State	Zip	Country
Milton	MA	02186	US

☐ Additional joint inventors are named on the _____ supplemental sheet(s) PTO/AIA/10 attached hereto.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Patent Application Fee Transmittal

Application Number:	15808201
Filing Date:	09-Nov-2017
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Filer:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
CLAIMS IN EXCESS OF 20	2202	3	50	150
INDEPENDENT CLAIMS IN EXCESS OF 3	2201	1	230	230
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Extension - 1 month with $0 paid	2251	1	100	100
Miscellaneous:				
SUBMISSION- INFORMATION DISCLOSURE STMT	2806	1	120	120
		Total in USD ($)		**600**

Electronic Acknowledgement Receipt

EFS ID:	35475257
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice/Lindsey Lovato
Filer Authorized By:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03
Receipt Date:	20-MAR-2019
Filing Date:	09-NOV-2017
Time Stamp:	13:37:14
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$600
RAM confirmation Number	032019INTEFSW13385600
Deposit Account	
Authorized User	

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		NRLA-001US03_Amendment.pdf	914939 21e0b97a4e23f53376b61612154199ba686787b1	yes	52
	Multipart Description/PDF files in .zip description				
	Document Description		**Start**	**End**	
	Applicant Arguments/Remarks Made in an Amendment		23	52	
	Claims		2	22	
	Amendment/Req. Reconsideration-After Non-Final Reject		1	1	
Warnings:					
Information:					
2	Information Disclosure Statement (IDS) Form (SB08)	NRLA001US03_SB08_03-19-2019.pdf	1054067 5587b617031db8700636ae42b100f141f40e6d33	no	4
Warnings:					
Information:					
3	Oath or Declaration filed	NRLA-001US03_DECS.pdf	3806662 bf28dc0df24c2534a73d59bc9b094a4ae508647c	no	9
Warnings:					
Information:					
4	Non Patent Literature	BrainInspiredComputing_Boston2010.pdf	21467 f8d229582ce1aca4440876f9b57c19a693e3e2e7	no	1
Warnings:					
Information:					
5	Non Patent Literature	Livitz_AdaptiveBrain-likeSystems.pdf	459911 a7f8963c9c7e1b65e499f57e2e7f39bcddcb2fb7	no	3

Warnings:					
Information:					
6	Non Patent Literature	SherbakovCogeyeYouTube.pdf	116888 6544fff61fc6b018d29663519012a4599a5fa8e6	no	1
Warnings:					
Information:					
7	Non Patent Literature	SherbakovCogeyeSystemDiagram.pdf	957468 e57f2887c3939a4c3034dc7028fe73f5768d2482	no	1
Warnings:					
Information:					
8	Non Patent Literature	Snider_IEEE_Computer_2011.pdf	1623491 6251a15257c12738f74a717e3c3988e7de85451a	no	9
Warnings:					
Information:					
9	Non Patent Literature	Versace_TedX_Fulbright_2014.pdf	4455341 aba040ae40d2f5ecd2790566960f76728498d775	no	30
Warnings:					
Information:					
10	Non Patent Literature	Salakhutdinov_Conference_on_Artifical_Intelligence_and_Statistics_2009.pdf	553762 284b18d7196f608448ca3d9496bf220b1dfffcf5	no	8
Warnings:					
Information:					
11	Fee Worksheet (SB06)	fee-info.pdf	35413 f3f0df7267f5fa9580d30f26f6878b194b4e7ddd	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	13999409	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Table of Contents

1 Status of the Claims and Summary of Amendments 25
2 Summary of Interview with Examiner 27
3 Prior and Concurrent Proceedings 28
4 Information Disclosure Statement 28
5 Rejections Under 35 U.S.C. § 112 28
6 Objection to Reissue Declaration 28
7 Rejections under 35 U.S.C. § 251 – Reissue Declaration 28
8 Rejections under 35 U.S.C. § 251 – Recapture 28
8.1 Claim 21 29
8.2 Claim 42 30
8.3 Claim 56 32
8.4 Claim 62 33
8.5 Claims 68 and 79 33
8.6 Claim 90 35
8.7 Claim 95 36
8.8 Claim 101 37
8.9 Claim 108 39
8.10 Claim 112 39
9 Rejections under 35 U.S.C. § 102 over Buck 40
9.1 U.S. Patent No. 7,219,085 B2 to Buck 40
9.2 Claim 1 45
9.3 Claim 3 45
9.4 Claim 11 46
9.5 Claim 12 46
9.6 Claim 21 47
9.7 Claim 42 47
9.8 Claim 56 47
9.9 Claim 62 47
9.10 Claim 68 47
9.11 Claim 79 48
9.12 Claim 90 48
9.13 Claim 95 48
9.14 Claim 117 49
9.15 Claim 128 49

10 Rejections under 35 U.S.C. § 102 over Burton 49
10.1 U.S. Pre-Grant Publication No. 2001/0010034 A1 to Burton 49
10.2 Claim 101 49
10.3 Claim 108 50
10.4 Claim 112 50
11 Rejections under 35 U.S.C. § 103 50

REMARKS

In response to the Office Action dated November 20, 2018, in the above-identified application, Applicant respectfully requests reconsideration. All of the issues raised in the Office Action have been carefully considered and are addressed in this response. The application as now presented is believed to be in allowable condition.

For brevity, Applicant is presenting remarks with respect to the rejections of the independent claims and certain representative dependent claims. This brevity should not be taken as a waiver of any issue with respect to any claim or acquiescence to a rejection.

1 Status of the Claims and Summary of Amendments

Upon entry of this response, claims 1–141 will be pending in this application, of which claims 1, 12, 21, 42, 56, 62, 68, 79, 90, 95, 101, 108, 112, 117, 128, and 137 are independent claims.

Claim 1 is being amended to read "the at least one graphics processing unit" instead of "the at least one graphic processing unit." Support for this change appears elsewhere in claim 1.

Claim 11 is being amended to read "the central processing unit" instead of "the central processor unit." Antecedent basis and support for "the central processing unit" appears in independent claim 1.

Claim 12 is being amended to provide proper antecedent basis for "the input data" and "the output data" in claims 12–14 and 16–19.

Claim 21 is being amended to recite "computations in the sequence of computations represent[] respective layers of neurons in an artificial neural network." Support for these amendments appears at least in column 5, lines 22–26 and 61–63 and column 9, lines 3–7 of U.S. Patent No. 9,189,828 ("the '828 patent").

Claim 55 is being amended to correct an informality ("from from" is being changed to "from").

Claim 56 is being amended to include claim 61.

Claim 61 is being canceled.

Claim 62 is being amended to include claim 67.

Claim 67 is being canceled.

Claim 68 is being amended to recite that the inputs are "inputs to the artificial neural network." Support for this amendment appears at least in column 8, lines 6–22 of the '828 patent.

Claim 79 is being amended to recite "a first memory partition" and "a second memory partition" instead of "at least one memory" and to recite "inputs to the artificial neural network." Support for these amendments appears at least in FIG. 2, which shows a texture memory bank 250 with four partitions, and in column 8, lines 6–22 of the '828 patent.

Claim 88, which depends from claim 78, is being amended to conform to claim 78 as amended.

Claim 90 is being amended to provide proper antecedent basis for "the output data" in claims 90 and 94.

Claim 95 is being amended to provide proper antecedent basis for "the input data" and "the output data" in claims 95 and 98–100.

Claim 101 is being amended to recite "a numerical simulation representing an artificial neural network," "a first computational cycle representing a first layer of neurons in the artificial neural network," "a second computational cycle representing a second layer of neurons in the artificial neural network," "first output data representing at least one output of the first layer of neurons," and "second output data representing at least one output of the second layer of neurons." Support for these amendments appears at least in the following sections of the '828 patent: column 1, lines 38–41; column 5, lines 22–26; column 5, lines 61–63; and column 9, lines 3–7.

Claim 108 is being amended in a fashion similar to claim 101.

Claim 112 is also being amended in a fashion similar to claim 101.

Claim 115, which depends from claim 112, is being amended to read "the controller" instead of "the accelerator controller" and to conform to claim 112 as amended. Antecedent basis for "the controller" appears in claim 112.

Claim 116, which depends from claim 112, is being amended to claim 112 as amended.

Claim 117 is being amended to read "the at least one processing unit" instead of "the at least one graphic processing unit." Support for this change appears elsewhere in claim 117.

Claim 127 is being amended to read "the central processing unit" instead of "the central processor unit." Antecedent basis and support for "the central processing unit" appears in independent claim 117.

Claim 128 is being amended to provide proper antecedent basis for "the input data" and "the output data" in claims 128–130 and 132–135.

Claims 137–141 are being added. Support for these claims appears at least in column 8, lines 6 – 12; column 9, lines 8–13; and column 12, lines 54–67 of the '828 patent. These new claims are patentable over the prior art of record at least because they require "a first processing unit to receive input data *in real time from a sensor*" and "a controller ... to store a second portion of the input data in the memory *while* the second processing unit generates the first output data." And they are not subject to recapture because they are materially narrower than the issued claims, at least for reciting "a first processing unit to receive input data *in real time from a sensor*."

2 Summary of Interview with Examiner

Applicant's representatives Christopher Max Colice and Andrew Baluch thank Examiner Wood and his conferees for granting and conducting an interview on February 8, 2019. During the interview, the participants discussed Applicant's technology, the prior art, and the recapture rejections. Examiner Wood also explained that several references listed in the Information Disclosure Statement were not considered because no copies could be located in the file.

In particular, the participants discussed Applicant's technology, claim 1, and the GPU-based training phase in columns 11 and 12 and FIG. 5 of U.S. Patent No. 7,219,085 B2 to Buck. Applicant's representatives explained that, in Applicant's technology, the sequence of computations could represent an artificial neural network and each computation could represent a layer of neurons. When implementing an artificial neural network, the controller of claim 1 would "transfer the at least a portion of the input data into the accelerator memory ... *during* performance of the sequence of computations [the artificial neural network] by the at least one graphics processing unit." In contrast, the cited section of Buck only discloses transferring data into GPU memory *before* the GPU performs any computations.

Examiner Wood then explained that each computation in claim 1 could be interpreted to represent an entire pass through Buck's training phase, in which case the sequence of computations would be many passes through Buck's training phase. Thus, a sequence of training or testing phases could involve transferring data to the GPU memory many times. But, as explained below, closer examination of Buck's columns 11 and 12 and FIG. 5 reveals that, even under this interpretation, Buck's training phase still doesn't include transferring data into the

GPU memory while the GPU is performing any computations. Instead, the transfer and computations occur in an alternating fashion as explained in Sections 9.1–9.3 starting on page 40 below.

3 Prior and Concurrent Proceedings

To Applicant's knowledge, there are no prior or concurrent judicial proceedings involving U.S. Patent No. 9,189,828.

4 Information Disclosure Statement

Applicant is enclosing an information disclosure statement listing the documents that were lined through in the Form 1449 attached to the Office Action. Copies of these documents are also enclosed per 37 C.F.R. § 1.98.

5 Rejections Under 35 U.S.C. § 112

Claims 1–20, 90–100, 115, and 117–136 were rejected as indefinite for lack of antecedent basis. Applicant respectfully submits that the amendments above provide proper antecedent basis so these rejections should be withdrawn.

6 Objection to Reissue Declaration

In page 3 of the Office Action, the reissue declaration was objected as defective because it failed to indicate the antecedent basis errors of claim 1 or how claim 21 is broadened. Applicant is enclosing revised reissue declaration to address this objection.

Applicant thanks Examiner Wood for his very helpful suggestions in revising the reissue declaration.

7 Rejections under 35 U.S.C. § 251 – Reissue Declaration

Claims 1–136 were rejected because the error statement in the reissue declaration was allegedly insufficient. Applicant submits that the revised reissue declaration overcomes this rejection.

8 Rejections under 35 U.S.C. § 251 – Recapture

Applicant respectfully requests withdrawal of the recapture rejections of claims 21–117 for at least the following reasons. As explained below, recapture does not apply to the rejected

claims for one or more of the following reasons: the rejected claims are materially narrower in certain respects than the claims issued in U.S. Patent No. 9,189,828 ("the '828 patent"), the rejected claims do not omit surrendered subject matter, or the rejected claims are directed to overlooked aspects.

8.1 Claim 21

Claim 21 has been rejected for recapture with respect to claim 1. Applying the three-step recapture test shows that recapture doesn't apply to claim 21 because (1) it includes the limitation ("the controller") added to define claim 1 over the art and (2) it is materially narrower than claim 1 in certain respects.

First, amended claim 21 is broader than claim 1 because it doesn't include all the elements of claim 1. For instance, amended claim 21 doesn't recite the "central processing unit," "main memory," or "accelerator" of claim 1.

Second, the broader aspects of claim 21 do not relate to subject matter surrendered during prosecution of claim 1. Per MPEP § 1412.02(II)(B)(1)(A), subject matter is surrendered when it is relied on to make the claims allowable over the prior art:

> If an original patent claim limitation now being omitted or broadened in the present reissue application was originally relied upon by applicant in the original application to make the claims allowable over the art, the omitted limitation relates to subject matter previously surrendered by applicant.

In this case, Applicant relied upon the "controller" to make claim 1 allowable over the prior art. Specifically, Applicant amended claim 1 to recite "a controller ... to transfer at least a portion of the output data from the accelerator memory to [a] main memory during performance of [a] sequence of computations by [a] graphic processing unit" and argued that this limitation made claim 1 allowable over the art.[1] (The "controller" in claim 21 is essentially the same as the "controller" in claim 1.) There is no recapture because claim 21 includes the entirety of the limitation ("the controller") added to define claim 1 over the art.[2]

---

[1] Amendment and Response filed March 10, 2015 in U.S. Application No. 14/147,015, page 6 ("Applicants respectfully submit that neither Chung nor Evans discloses or suggests "a controller ... to transfer at least a portion of the output data from the accelerator memory to [a] main memory during performance of [a] sequence of computations by [a] graphic processing unit" as recited in claim 1 as amended.").

[2] MPEP § 1412.02(III)(A)(2) ("If the reissue claims are narrower in scope than the claims canceled from the original application by inclusion of the entirety of the limitation added to define the original application claims over the art, there will be no recapture, even if the reissue claims are broader than the canceled claims in some other aspect (i.e., an aspect not related to the surrender made in the original application).").

Third, even if the broader aspects of claim 21 could be considered to relate to surrendered subject matter, amended claim 21 is materially narrower than claim 1 in aspects related to the limitations that made claim 1 allowable over the prior art. Per MPEP § 1412.02(II), "if the reissue claims were materially narrowed in aspects related to the surrendered subject matter, there is no recapture...." More specifically, claim 21 recites "a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data, the sequence of computations representing an artificial neural network, computations in the sequence of computations representing respective layers of neurons in an artificial neural network." This is a modified version of the surrender-generating limitation of claim 1. The modification—limiting the "sequence of computations" to represent "an artificial neural network," with computations representing respective layers of neurons—makes claim 21 narrower than claim 1 in a meaningful way. That's because an artificial neural network is a particular type of the sequence of computations claimed broadly in claim 1:

> The elements of computation (computational elements) are characterized by a single output variable. Such computational elements often include one or more equations. Computational elements are same or similar within a large population and are computed in parallel. *An example of such a population is a layer of neurons in an artificial neural network (ANN), where all neurons are described by the same equation.*[3]

Because this materially narrower limitation is related to the limitations that made claim 1 allowable over the prior art, claim 21 avoids recapture.

8.2 Claim 42

Recapture does not apply to amended claim 42 with respect to claim 12 because amended claim 42 is directed to a "separate species or embodiment" than claim 12.[4] Claim 42 is directed to a separate species or embodiment because, as pointed out in page 6 of the Office Action, it includes hardly any of the language of claim 12. In fact, claims 12 and 42 are so dissimilar that they cannot be compared word-for-word.

---

[3] '828 patent, 5:20–25.

[4] *See* MPEP 1412.01(II) ("Claims to separate inventions/embodiments/species that were disclosed but never covered by the claims in the original application prosecution are claims to overlooked aspects. In other words, the reissue claims are drawn to a separate invention or separate species or embodiment that was not covered by a claim (e.g., a generic claim) at any point during the prosecution of the original application.... Recapture analysis would not need to be continued for claims drawn to overlooked aspects.").

If anything, amended claim 42 is probably more properly tested for recapture with respect to claim 1 because claim 42 recites limitations more similar to those in claim 1 than claim 12. Nevertheless, applying the three-step recapture test with respect to claim 1 shows that recapture doesn't apply to claim 42 with respect to claim 1 because (1) the broader aspects of claim 42 don't relate to surrendered subject matter and (2) claim 42 is materially narrower than claim 1 in aspects that relate to the limitations that made claim 1 allowable over the art:

Claim 1	Claim 42
1. A computer system, comprising:	42. A method of performing a sequence of computations representing an artificial neural network, the method comprising:
a central processing unit to receive input data;	receiving, at a central processing unit (CPU), first input data acquired from an external system;
main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;	
an accelerator, operably coupled to the central processing unit and the [first] main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:	
at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations yielding intermediate results; and	performing, by a graphics processing unit (GPU), a first computation in the sequence of computations on the first input data to generate first output data, computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of a computation in the sequence of computations representing an output of a neuron in the artificial neural network;
accelerator memory, operably coupled to the at least one graphics processing unit, to store the results of the sequence of computations; and	storing, in a memory operably coupled to the GPU, the first input data and the first output data; and
a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to transfer the at least a portion of the input data into the accelerator	transferring second input data into the memory during performance of the sequence of computations by the GPU.

memory and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one graphics processing unit.	

First, amended claim 42 is broader than claim 1 because it doesn't recite the main memory or accelerator of claim 1.

Second, the broader aspects of amended claim 42 do not relate to subject matter surrendered during prosecution of claim 1. As explained above with respect to claim 21, claim 1 was allowed because it was amended to recite "a controller ... to transfer at least a portion of the output data from the accelerator memory to [a] main memory during performance of [a] sequence of computations by [a] graphic processing unit."[5] Analogous method limitations appear in claim 42.

Third, amended claim 42 is materially narrower than claim 1 in an aspect related to the limitations that made claim 1 allowable over the art: it requires the sequence of computations occurring during the transfer to "represent[] respective layers of neurons in the artificial neural network" and "an output of a computation in the sequence of computations representing an output of a neuron in the artificial neural network." As explained above with respect to claim 21, these limitations restrict amended claim 42 to specific types of computations rather than the more generic computations recited in claim 1.

8.3 Claim 56

Like claim 21, claim 56 has been rejected for recapture with respect to claim 1. Claim 21 is broader than claim 1 because it doesn't include the "central processing unit," "main memory," or "accelerator" of claim 1. But like claim 21, claim 56 avoids recapture at step 2 because claim 56 recites the "controller" of claim 1, which was the only subject matter surrendered during prosecution of claim 1. And like claim 21, claim 56 further avoids recapture at step 3 because claim 56 is materially narrower than claim 1 in that it recites "a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data, a computation

[5] Amendment and Response filed March 10, 2015 in U.S. Application No. 14/147,015, page 6 ("Applicants respectfully submit that neither Chung nor Evans discloses or suggests "a controller ... to transfer at least a portion of the output data from the accelerator memory to [a] main memory during performance of [a] sequence of computations by [a] graphic processing unit" as recited in claim 1 as amended.").

in the sequence of computations comprising a plurality of computational elements, *the plurality of computational elements representing a layer of neurons in an artificial neural network, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network*."

8.4 Claim 62

Like claim 42, claim 62 has been rejected for recapture with respect to claim 12, but is more properly compared to claim 1. Claim 62 is broader in some respects than claim 1 but avoids recapture at step 3 of the recapture test because claim 62 is materially narrower than claim 1 in a manner related to the allowable limitations of claim 1. In particular, claim 62 recites "performing, with a graphics processing unit (GPU), a sequence of computations on the input data to generate output data, a computation in the sequence of computations comprising a plurality of computational elements, *the plurality of computational elements representing a layer of neurons in an artificial neural network, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network*." This limitation makes claim 62 materially narrower than claim 1.

8.5 Claims 68 and 79

Recapture does not apply to claims 68 and 79 with respect to the '828 patent because claims 68 and 79 are directed to a "separate species or embodiment" than the claims of the '828 patent.[6]

Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU accelerator and an architecture for the GPU accelerator, respectively. Claims 68 and 79, on the other hand, are directed to the embodiment of FIG. 3, which shows top-level control of the GPU accelerator system. Claims 68 and 79 recite "a user interaction stream" and "a computational stream," neither of which appear in the claims of the '828 patent, FIGS. 1 or 2, or the descriptions of FIGS. 1 and 2.

[6] *See* MPEP 1412.01(II) ("Claims to separate inventions/embodiments/species that were disclosed but never covered by the claims in the original application prosecution are claims to overlooked aspects. In other words, the reissue claims are drawn to a separate invention or separate species or embodiment that was not covered by a claim (e.g., a generic claim) at any point during the prosecution of the original application.... Recapture analysis would not need to be continued for claims drawn to overlooked aspects.").

Although the embodiment of FIGS. 1 and 2 are related to the embodiment of FIG. 3, these embodiments are distinct inventions. According to MPEP § 802.01,

> [R]elated inventions are distinct if the inventions as claimed are not connected in at least one of design, operation, or effect (e.g., can be made by, or used in, a materially different process) and wherein at least one invention is PATENTABLE (novel and nonobvious) OVER THE OTHER (though they may each be unpatentable over the prior art).

In this case, although FIGS. 1, 2, and 3 relate to a GPU accelerator, FIGS. 1 and 2 illustrate the architecture of a computer system that includes a GPU accelerator and the GPU accelerator itself, respectively. Conversely, FIG. 3 shows top-level control of the GPU accelerator system. Although this top-level control can be used with the hardware in FIGS. 1 and 2, it can be used with other hardware as well. In addition, this top-level control involves execution streams and the interaction between execution streams in the CPU and GPU in order to control computation and process numerical techniques. These features make the embodiment of FIG. 3 patentably distinct over the embodiments of FIGS. 1 and 2.

Assuming *arguendo* recapture did apply to claims 68 and 79, it would apply with respect to ***claim 9*** of the U.S. Patent No. 8,468,867 ("the '867 patent"), which is the parent of the '828 patent, because claims 68 and 79 are more similar to claim 9 of the '867 patent than to any other issued claim in the '828 patent or the '867 patent. All of these claims recite a computational stream, for example, and are supported by FIGS. 3 and 4 and the accompanying description.

Claim 9 of the '867 patent recites:

> 9. An accelerator system for a computer system performing a numerical simulation, the accelerator system comprising:
>
> at least one graphics processing unit to generate output data by performing at least one computation during a first computational cycle of the numerical simulation;
>
> an accelerator memory, operably coupled to the at least one graphics processing unit, to store data used to perform the at least one computation; and
>
> an accelerator controller, operably coupled to the accelerator memory and the at least one graphics processing unit, to execute:
>
> (i) a computational stream controlling performance of the at least one computation by the at least one graphics processing unit;
>
> (ii) an output stream controlling transfer of the output data from the at least one graphics processing unit to the accelerator memory during the first computational cycle; and
>
> (iii) an input stream controlling transfer of input data to the accelerator memory for use by the at least one graphics processing unit during a second computational cycle of the numerical simulation.

Nevertheless, applying the three-step recapture test shows that recapture doesn't apply to claims 68 and 79 with respect to ***claim 9*** either:

First, claims 68 and 79 are broader than claim 9 because they don't include all the elements of claim 9 (for instance, neither claim 68 nor 79 recite "an output stream" or "an input stream" as in claim 9).

Second, the broader aspects of claims 68 and 79 do relate to subject matter surrendered during prosecution of claim 9, namely, the accelerator controller and the computational, output, and input streams.[7]

But third, recapture does not apply at step 3 because claims 68 and 79 are materially narrower than claim 9 in at least three respects: they each recite (1) "computations representing an artificial neural network"; (2) "a user interaction stream" that controls "transfer of inputs to the artificial neural network to the first memory partition and the second memory partition"; and (3) "a computational stream" that controls "data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network." An artificial neural network is more specific than the computational cycles that contributed to the patentability of claim 9 over the art. Likewise, the first and second memory partitions in claims 68 and 79 are more specific than the accelerator memory in added portions of claim 9. And the functions of the user interaction and computational streams in claims 68 and 79 are more specific than the functions of the computational, input, and output streams of claim 9. All of these more specific features make claims 68 and 79 materially narrower than claim 9 of the '867 patent.

8.6 Claim 90

Claim 90 was rejected for recapture with respect to claim 1. However, claim 90 is more properly compared to claim 12 (italics show the additional limitations in claim 90):

Claim 12	Claim 90

[7] Amendment filed on July 31, 2013 in U.S. Application No. 11/860,254 at page 14 ("Applicants respectfully submit that Chung and Evans, alone or in combination, fail to disclose or suggest all of the elements recited in claim 10 [issued claim 9 in the '867 patent]. More specifically, nowhere does Chung or Evans disclose or suggest an accelerator controller that executes a computational stream to control performance of a first computation, an output stream to transfer data at the end of the first computational cycle, and an input stream to control the transfer of input data into the accelerator memory. Therefore, no combination of Chung and Evans yields an accelerator system that includes the accelerator controller of claim 10 as amended.").

12. A method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory, the method comprising:	90. A method of performing a sequence of computations with a processing unit, the method comprising:
	receiving input data;
(A) performing, by the GPU, the sequence of computations on a first portion of input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results;	performing, by the processing unit, the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, *the first portion of the output data representing an output of an artificial neural network;*
(B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; and	in parallel with performing the sequence of computations by the processing unit, transferring a second portion of the input data *from a first memory partition to a second memory partition; and*
(C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus.	in parallel with performing the sequence of computations by the processing unit, transferring a second portion of the output data *from the second memory partition to the first memory partition.*

Although claim 90 has a shorter preamble than claim 12, it is narrower in every other respect. Per MPEP § 1412.02(II)(A), "[i]f the reissue claim is not broadened in any respect as compared to the patent claims, the analysis ends; there is no recapture." Therefore, recapture does not apply to claim 90 per the first step of the three-step recapture analysis.

8.7 Claim 95

Claim 95 is an apparatus claim analogous to claim 90 and has been rejected for recapture with respect to claim 1. Like claim 90, claim 95 is best compared with claim 12, not claim 1, and

is narrower in just about every respect than claim 12. Therefore, recapture does not apply to claim 95.

8.8 Claim 101

Claim 101 was rejected for recapture with respect to claim 1. However, claim 101 is more properly compared to claim 16 of the '867 patent (italics show the additional limitations in claim 90):

Claim 16, '867 patent	Claim 101
16. A method for performing a numerical simulation on input data in a computer system including a central processing unit and an accelerator, the method comprising:	101. A method for performing a numerical simulation representing an artificial neural network over a plurality of computational cycles including at least a first computational cycle representing a first layer of neurons in the artificial neural network and a second computational cycle representing a second layer of neurons in the artificial neural network, the method comprising:
receiving, by an accelerator, first input data from the central processing unit;	
transferring, by an accelerator controller, the first input data into a first partition, referenced by first pointer, of an accelerator memory before a first computational cycle of the numerical simulation;	storing, *in a first partition of a memory*, first output data during the second computational cycle, the first output data generated during the first computational cycle and being referenced by a first pointer, *the first output data representing at least one output of the first layer of neurons;*
performing, by at least one graphics processing unit during the first computational cycle, at least one calculation on the first portion of the input data as to generate first output data;	performing, by a processing unit operably coupled to the memory, at least one calculation on the first output data during the second computational cycle, the at least one calculation yielding second output data, *the second output data representing at least one output of the second layer of neurons;*
storing, by the accelerator controller, the first output data into a second partition, referenced by a second pointer, of the accelerator memory; and	storing, *in a second partition of the memory*, the second output data, the second output data being referenced by a second pointer; and

swapping the first pointer with the second pointer at the end of the first computational cycle, such that the first output data becomes an input for a second computational cycle of the numerical simulation.	swapping the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data *becomes an input for a third computational cycle of the plurality of computational cycles.*

Recapture does not apply to claim 101 with respect to claim 16 of the '867 patent because claim 101 is materially narrower than claim 16 with respect to subject matter that Applicant used to distinguish claim 16 from the prior art.[8]

Under the first step of the three-step recapture analysis, claim 101 is broader in some respects and narrower in other respects than claim 16.

Under the second step, the broader aspects of claim 101 relate to subject matter than Applicant relied on to distinguish claim 16 from the prior art. This subject matter includes first input data in a first partition referenced with a first pointer, first output data in a second partition referenced by a second pointer, and swapping the first and second pointers.[9]

Under the third step, however, claim 101 is materially narrower than claim 16 because it includes limitations that narrow the subject matter that Applicant used to distinguish claim 16 from the prior art. These narrower limitations include "first output data representing at least one output of the first layer of neurons," "a first memory partition" and "a second memory partition" instead of "an accelerator memory," "second output data representing at least one output of the second layer of neurons," and "swapping the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles." Thus, recapture does not apply to claim 101.

[8] *See* MPEP § 1412.02(II) ("[I]f the reissue claims were materially narrowed in aspects related to the surrendered subject matter, there is no recapture....").

[9] Amendment filed on July 31, 2013 in U.S. Application No. 11/860,254 at page 13 ("Like claim 1 as amended, independent claim 17 [issued as claim 16] as amended recites an accelerator controller that transfers 'the first input data into a first partition, referenced by first pointer, of an accelerator memory before a first computational cycle of the numerical simulation' and that loads first output data generated during the first computational cycle 'into a second partition, referenced by a second pointer, of the accelerator memory.' And Claim 17 as amended also recites 'swapping the first pointer with the second pointer at the end of the first computational cycle, such that the first output data becomes an input for a second computational cycle of the numerical simulation.' Thus, like the subject matter of claim 1 as amended, the subject matter of claim 17 as amended is patentably distinct from Chung and Evans.")

8.9 Claim 108

Claim 108 has been rejected for recapture with respect to claim 12. Applying the three-step recapture test shows that recapture doesn't apply to amended claim 21 with respect to claim 12 because the broader aspects of claim 108 don't relate to surrendered subject matter and because claim 108 is materially narrower than claim 12:

First, amended claim 108 is broader than claim 12 because it doesn't include all the elements of claim 12.

Second, the broader aspects of amended claim 108 do not relate to subject matter surrendered during prosecution of claim 12 because no subject matter was surrendered in claim 12 during prosecution. This alone means that claim 108 is not subject to recapture at step 2.

Third, even if the broader aspects of amended claim 108 do relate to surrendered subject matter, claim 108 is materially narrower than claim 12 because it includes several limitations that don't appear in claim 12. These limitations are related to the inventive aspects of claim 12 and include "a first computational cycle ... representing a first layer of neurons in the artificial neural network" and " a second computational cycle ... representing a second layer of neurons in the artificial neural network." Thus, the recapture test fails at step 3.

8.10 Claim 112

Claim 112 was rejected for recapture with respect to claim 1. However, claim 112 is more properly compared to claim 1 of the '867 patent because each of these claims recites a controller that swap pointers, among other things. Again, applying the three-step recapture analysis includes shows that recapture does not apply because claim 112 includes the limitation ("the controller") added to define claim 1 of the '867 patent over the art and is materially narrower than claim 1 of the '867 patent.

First, claim 112 is broader because it does not include all the limitations of claim 1 of the '867 patent.

Second, the broader aspects of claim 112 relate to subject matter surrendered during prosecution of claim 1 of the '867 patent. Claim 1 was allowed over the prior art because it includes an accelerator controller.[10]

[10] Amendment filed on July 31, 2013 in U.S. Application No. 11/860,254 at pages 12 and 13 ("More specifically, nowhere in the reference does Evans disclose or suggest an accelerator controller, operably coupled to the accelerator memory and the central processing unit, to transfer the at least the portion of the input data into the accelerator memory before the first computational cycle, to transfer the first output data from the accelerator memory to the main memory during the second computational cycle, to direct the second output data into the second

Third, claim 112 is materially narrower than claim 1 because it specifies computing layers of neurons instead of more generic computational cycles and that "the first output data becomes an input for a second layer of neurons of the artificial neural network." These limitations are related to the limitations that made claim 1 allowable over the art and restrict claim 112 to specific applications instead of more general computation. Thus, claim 112 avoids recapture under step 3.

9 Rejections under 35 U.S.C. § 102 over Buck

Claims 1–136 have been rejected as anticipated by Buck. Applicant respectfully requests withdrawal of these rejections for at least the following reasons.

9.1 U.S. Patent No. 7,219,085 B2 to Buck

Buck discloses using a GPU to accelerate and optimize training and testing of a two-layer neural network. In Buck's system, a CPU transfers input data, programs, and parameters to the GPU before the GPU performs any calculations. These programs can instruct the GPU to train or test parameters for a neural network. The GPU executes the programs on the input data using the parameters, then transfers output data to the CPU. If desired, the testing can be repeated for new input data.

Buck's FIG. 2 (reproduced below) illustrates a computer 110 that can be used to carry out this acceleration and optimization. This computer 110 includes CPU 120 coupled to a CPU memory 130. The CPU 120 is also coupled, via an accelerated bus 210, to a video interface 190, which includes a GPU 220, GPU memory 230, and machine learning GPU execution engine 200. This machine learning GPU execution engine 200 includes a training module 240 and a testing module 250.

partition at the end of the second computational cycle, and to swap the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles as recited in independent clam 1.")

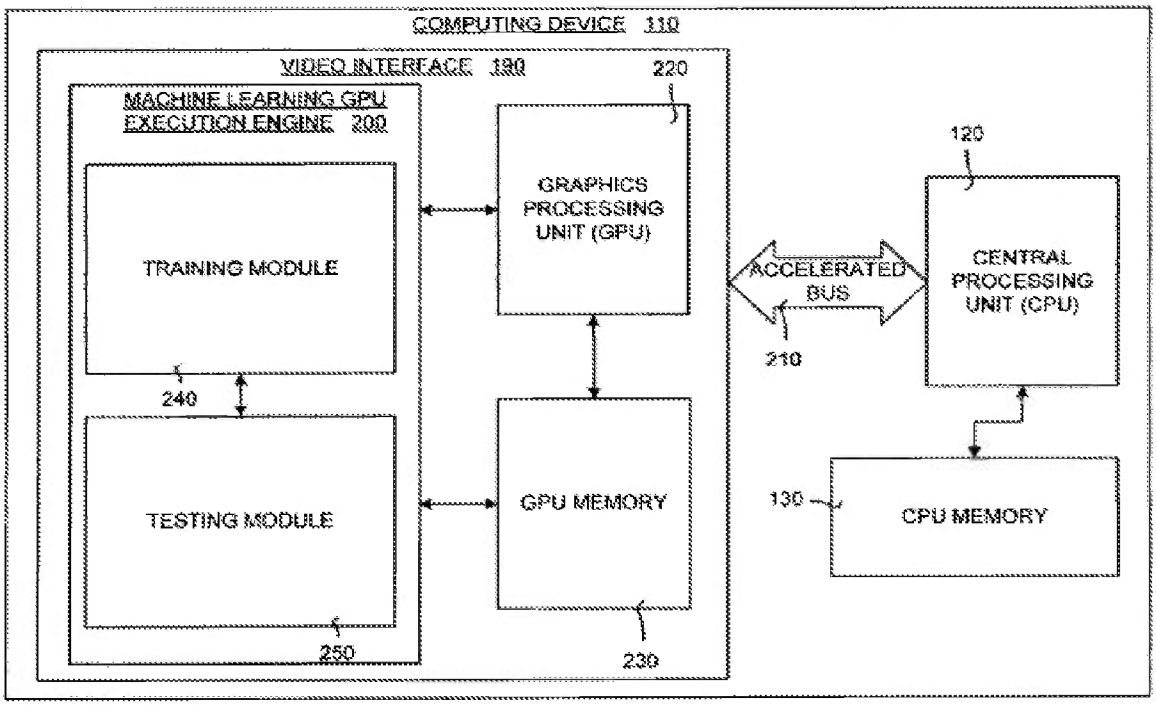


FIG. 2

Buck's computer 110 performs machine learning in two phases: (1) a training phase, carried out by the training module 240, in which the parameters of the algorithm are adjusted using training data; and (2) a testing phase, carried out by the testing module 250, in which machine learning parameters can be computed as a function of input data and training parameters.[11] One example of this machine learning is training a two-layer neural network using backpropagation, where the machine learning parameters are weights for the layers in the neural network.[12]

Buck's FIG. 5 (reproduced below) illustrates the training phase.[13] In the first step of the training phase, the CPU loads shaders programs (P), trainable parameters (W), and learning parameters (L) into the GPU memory (arrow 505 in FIG. 5).[14] For a neural network, the learning parameters are weights for the layers in the neural network.[15] The shaders programs are used for

[11] Buck, 9:27–33.
[12] Buck, 21:35–23:32; *see also id.* at 11:42–45 and 12:43–46.
[13] Buck, 4:58–60 and 10:36–12:39.
[14] Buck, 10:36–50.
[15] Buck, 10:45–47 ("In a preferred implementation, the learning parameters (L) are the weights of each of the neural network layers and the threshold of each unit.").

updating the weights and propagating the weights forward and backward through the layers in the neural network.[16]

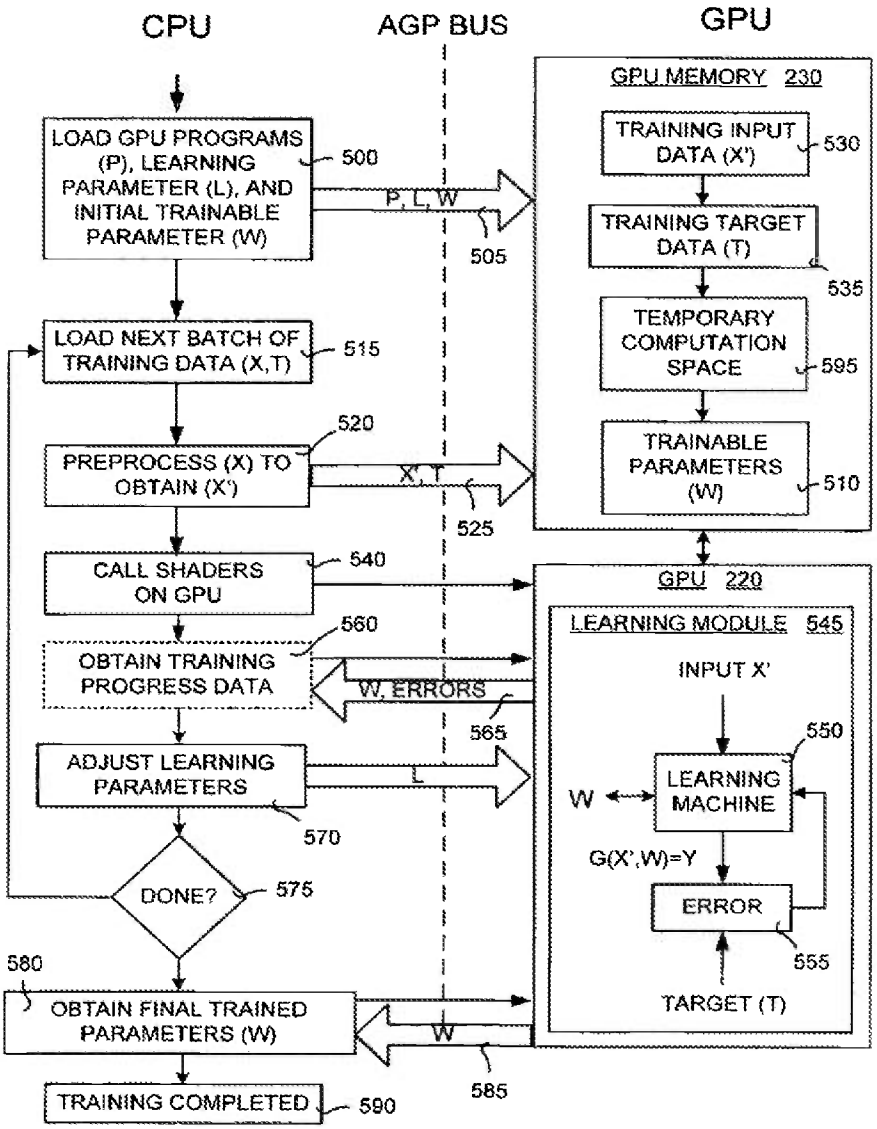


FIG. 5

Next, the CPU loads batches of training data (X) into the GPU memory (arrow 525 in FIG. 5).[17] For training a neural network, the input data is a training set, such as the MNIST set of images of handwritten digits,[18] and only needs to be loaded once, before training begins, because

[16] Buck, 10:48–50.
[17] Buck, 10:51–11:11.
[18] Buck, 21:54–64.

it doesn't change during training. The CPU doesn't load anything into the GPU or GPU memory when the GPU is carrying out the training.

"Once the training data (X') has been loaded on the GPU, the CPU instructs the GPU to run the various shaders needed to process the machine learning technique (box 540 [in FIG. 5])."[19] For a neural network, this means simulating every layer of neurons in the neural network.[20] At the end of the training phase, the GPU sends the training parameters back to the CPU (arrow 585 in FIG. 5).[21] "Training typically takes a long time, but it can be done once and for all."[22]

Buck is silent on transferring data from the CPU to the GPU while the GPU is executing the shaders.

In the testing phase, shown in Buck's FIG. 6 (reproduced below), the CPU loads shader programs (P), trainable parameters (W), and input data (X') onto the GPU memory (arrows 605 and 625 in FIG. 5).[23] Once the programs, parameters, and data have been loaded into the GPU memory, the CPU instructs the GPU to use the shader programs.[24] The GPU uses the shader programs along with the trained function to compute the output.[25] The GPU transfers results (Y) of testing back to the CPU at the end of testing (arrow 665 in FIG. 6).[26] If there is more data to test, the testing phase repeats.[27]

Buck is silent on transferring data from the CPU to the GPU while the GPU is executing the shader programs in training phase.

[19] Buck, 11:29–31.
[20] *See* Buck, 21:35–23:32.
[21] Buck, 12:24–26.
[22] Buck, 8:63–65.
[23] Buck, 12:40–63.
[24] Buck, 13:4–6.
[25] Buck, 13:6–13.
[26] Buck, 13:13–17.
[27] Buck, 13:19–22.

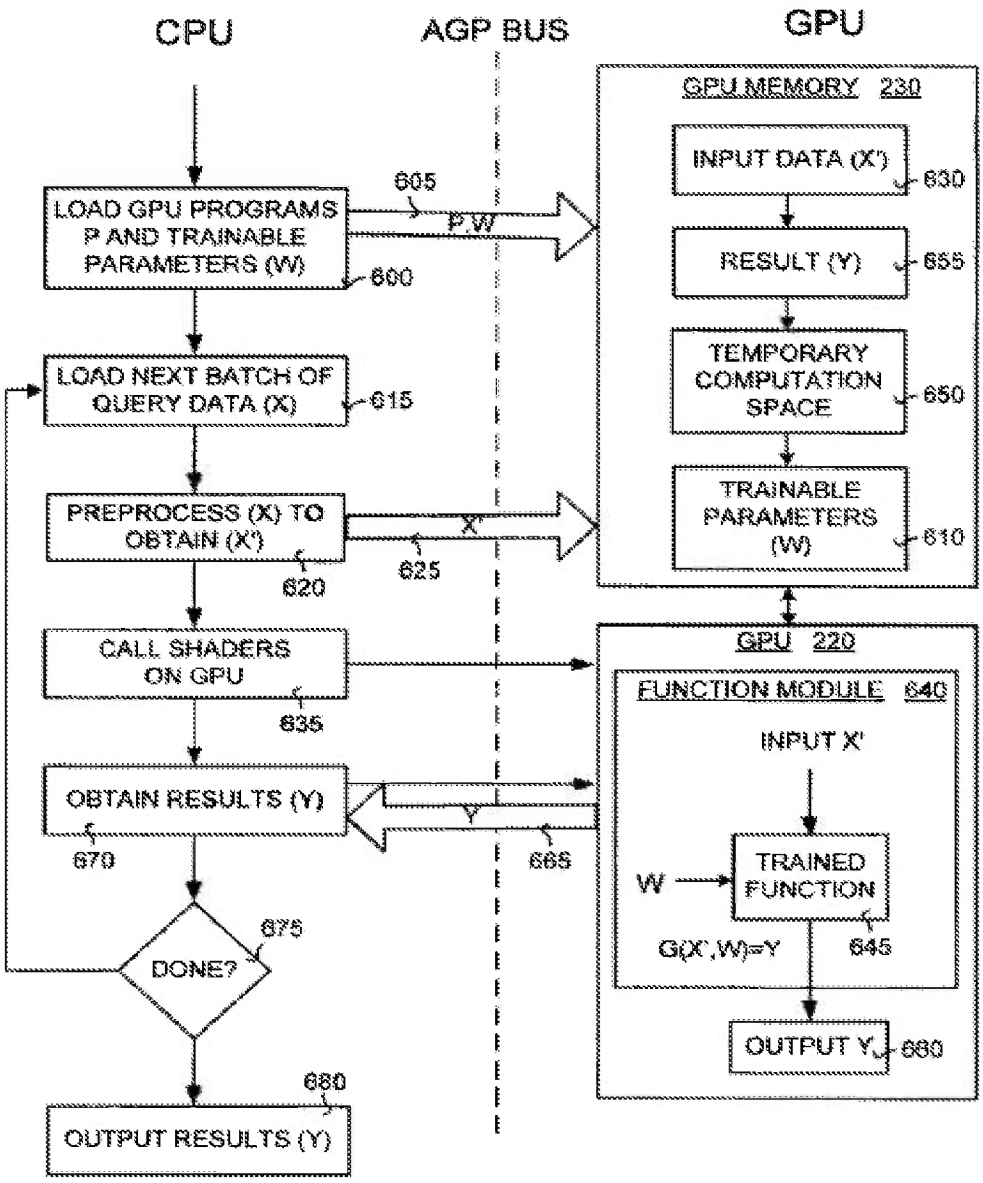


FIG. 6

Buck also gives a working example of using GPU to update the weights of a two-layer neural network using backpropagation.[28] In this example, each pass through the training phase comprises computing the outputs of both layers of the neural network (forward propagation), computing the error at the output of the second layer, updating the weights for the second layer,

[28] Buck, 21:35–23:32.

backpropagating the error from the second layer to the first layer, and updating the weights of the first layer.[29] At the end of each pass, the GPU can copy the updated weights to CPU.[30]

Again, Buck is silent on transferring data to the GPU while the GPU is training the neural network.

9.2 Claim 1

Claim 1 as amended recites "a controller ... to transfer the at least a portion of the input data into the accelerator memory ... during performance of the sequence of computations by the at least one graphics processing unit." As noted above, Buck is silent on transferring data to a GPU or GPU memory while the GPU is performing a sequence of computations (executing shaders programs). Although Buck discloses transferring shaders programs, learning parameters, trainable parameters, and input data from the CPU to the GPU memory in both training and testing phases, *all of these transfers occur before the GPU begins executing the shaders programs*. Thus, they cannot be used to provide the controller of claim 1 as amended.

This holds true under the mapping of Buck to claim 1 proposed during the interview by Examiner Wood. In this mapping, each pass through the training or testing phase counts as a "computation" in a "sequence" of multiple passes through the training or testing phase. Thus, a sequence of training or testing phases could involve transferring data to the GPU many times. But even when the GPU repeats the training or testing phase, it does not execute the shaders programs while the CPU is transferring data to the GPU. Thus, even under this mapping, Buck still doesn't disclose "a controller ... to transfer the at least a portion of the input data into the accelerator memory ... *during* performance of the sequence of computations by the at least one graphics processing unit" as in claim 1.

9.3 Claim 3

Claim 3 recites that "the central processing unit is configured to receive the input data at a first rate; and the at least one graphics processing unit is configured to perform the sequence of computations at a second rate different than the first rate." It was rejected over three sections of Buck:

[29] Buck, 22:1–23:9.
[30] Buck, 23:7–9 ("8. copy (matrix or vector) from CPU memory to GPU memory, or vice-versa (this is directly supported by DirectX, and mentioned only for completeness).").

(1) column 2, lines 55–67, which states just that "the GPU is typically more powerful than the CPU";

(2) column 9, lines 10–14, which states that "[a] GPU implementation increases the machine learning processing speed, both for training and testing, by up to an order of magnitude over CPU-only implementations"; and

(3) column 14, lines 40–44, which states "the training and use of these algorithms run an order of magnitude faster on a GPU than on a CPU."

None these sections say anything about the rate at which the CPU receives input data. In fact, Buck is silent on the rate at which the CPU receives input data. Thus, Buck does not disclose a CPU that "is configured to ***receive*** the input data at a first rate" and a GPU that "is configured to perform the sequence of computations at a second rate different than the first rate" as in claim 3.

9.4 Claim 11

Claim 11 recites "at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time." It was also rejected over three sections of Buck:

(1) column 7, line 51 to column 8, line 2, which states merely "[o]ther input devices ... may include a microphone ... or the like";

(2) column 9, lines 33–36, which mentions "speech recognition"; and

(3) column 23, lines 26–30, which says "there is no need to read back data from the graphics memory to the CPU unless the programmer wishes to monitor progress."

None these sections say anything about acquiring input data *in real time* with a video camera, microphone or cell recording electrode. In fact, Buck is silent on acquiring input data in real time. Thus, Buck does not disclose the system of claim 11.

9.5 Claim 12

Buck does not disclose "in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus." In contrast, Buck discloses transferring input data to the GPU and performing computations by the GPU in a sequential or interleaved fashion: first, the CPU transfers the input data to the GPU memory; then, once the input data is loaded, the GPU executes the shaders programs. These two steps do not occur at the same time—they occur one

after the other. Thus, Buck does not disclose transferring input data *in parallel* with performance of the sequence of computations by the GPU.

9.6 Claim 21

Like claim 1, claim 21 recites "a controller ... to transfer the input data into the memory and to transfer at least a portion of the output data from the memory during performance of the sequence of computations by the GPU." And like claim 1, claim 21 is patentably distinct from Buck because Buck is silent on transferring input data into memory during performance of a sequence of computations by the GPU.

9.7 Claim 42

Claim 42 recites, among other things, "transferring second input data into the memory [operably coupled to the GPU] during performance of the sequence of computations by the GPU." As explained above, Buck is silent on transferring input data into GPU memory while the GPU is performing computations.

9.8 Claim 56

Claim 56 recites "a controller ... to transfer at least a portion of the input data into the memory in parallel with performance of the sequence of computations by the GPU." As explained above, Buck is silent on transferring data to the GPU or GPU memory *in parallel with* performing computations on the GPU. Rather, Buck discloses transferring input data, learning parameters, trainable parameters, and shaders programs to the GPU in *a sequential or interleaved fashion* with the GPU's execution of the shaders programs.

9.9 Claim 62

Like claim 1, claim 62 recites "transferring the input data into the memory and at least a portion of the output data from the memory during performance of the sequence of computations by the GPU." And like claim 1, claim 62 is patentably distinct from Buck because Buck is silent on transferring input data into memory while the GPU is performing computations.

9.10 Claim 68

Claim 68 as amended recites "a user interaction stream" that "control[s] transfer of inputs to the artificial neural network to the first memory partition and the second memory partition."

Buck fails to provide the user interaction stream of amended claim 68. Page 36 of the Office Action maps the first and second memory partitions to the CPU and GPU memories, respectively, and suggests that data transfer to the CPU and GPU memories provides the user interaction stream of claim 68.[31] More specifically, the Office Action implies that the transfer of computation results from the GPU to CPU corresponds to the transfer of inputs to the first memory partition. But the data transferred from the GPU to CPU is not an input to an artificial neural network. Rather, it is the output of the training or testing phase. Thus, it cannot be used to provide the user interaction stream of claim 68.

9.11 Claim 79

Like amended claim 68, amended claim 79 recites "a user interaction stream" that "control[s] transfer of inputs to the artificial neural network to the first memory partition and the second memory partition." Amended claim 79 is therefore patentably distinct from Buck for at least the same reasons as amended claim 68.

9.12 Claim 90

Claim 90 recites, among other things, "in parallel with performing the sequence of computations by the processing unit, transferring a second portion of the input data from a first memory partition to a second memory partition." Claim 90 is therefore patentably distinct from Buck for at least the reasons given above with respect to claim 12.

9.13 Claim 95

Buck does not disclose or suggest "a first memory partition ... configured to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit" as in claim 95.

[31] At page 36, the Office Action cites two sections of Buck as providing the user interaction stream: column 9, lines 35–36 ("The results of the computation are transferred from the machine learning GPU execution engine 200 to the CPU 120 for use in interactive use applications, such as speech recognition.") and column 10, lines 41–43 ("Data transfer from the CPU to the GPU is relatively expensive, so the CPU pre-loads as much as possible onto the GPU before entering a training loop.").

9.14 Claim 117

Buck does not disclose "a controller ... transfer the at least a portion of the input data into the accelerator memory ... during performance of the sequence of computations by the at least one graphic processing unit."

9.15 Claim 128

Buck does not teach "in parallel with performing the sequence of computations by the processing unit in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus" as in claim 128.

10 Rejections under 35 U.S.C. § 102 over Burton

10.1 U.S. Pre-Grant Publication No. 2001/0010034 A1 to Burton

Burton discloses a method of simulating pipelined hardware, such as microprocessors. The hardware is simulated as circuit elements that read their inputs from a first data storage area and write their outputs to a second data storage area during a first clock cycle.[32] "At the end of each simulated common clock signal cycle, pointers to the two data storage areas are swapped such that the output values of the preceding cycle become the input values for the next stage in the following cycle."[33] Burton is silent on artificial neural networks.

10.2 Claim 101

Claim 101 is being amended to recite "a numerical simulation representing an artificial neural network," "a first computational cycle representing a first layer of neurons in the artificial neural network," "a second computational cycle representing a second layer of neurons in the artificial neural network," "first output data representing at least one output of the first layer of neurons," and "second output data representing at least one output of the second layer of neurons." Burton is silent on neural networks and therefore does not provide all the elements of claim 101 as amended.

[32] Burton, Abstract.
[33] Burton, paragraph 50.

10.3 Claim 108

Claim 108 is being amended to recite "a numerical simulation representing an artificial neural network," "a first computational cycle ... representing a first layer of neurons in the artificial neural network," and "a second computational cycle ... representing a second layer of neurons in the artificial neural network." Burton is silent on neural networks and therefore does not provide all the elements of claim 108 as amended.

10.4 Claim 112

Claim 112 is being amended to recite "[a] system for executing an artificial neural network" and a processing unit that "perform[s], during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data, the first output data representing an output of at least one neuron in the first layer of neurons." Burton is silent on neural networks and therefore does not provide all the elements of claim 112 as amended.

11 Rejections under 35 U.S.C. § 103

Claims 8, 9, 17, 18, 44–46, 96, 124, 125, 133, and 134 were rejected as obvious over Buck in view of U.S. Patent No. 7,525,547 B1 to Diard.

Claims 73, 78, 84, and 89 were rejected as obvious over Buck in view of U.S. Pre-Grant Publication No. 2006/0129506 A1 to Edelman *et al.*

Claim 105 was rejected as obvious over Burton in view of U.S. Patent No. 5,142,665 to Bigus.

Claims 110 and 113–116 were rejected as obvious over Burton in view of Buck.

Each of these claims depends from a patentable base claim, so Applicant believes that it is unnecessary at this time to argue the patentability of each of these claims individually. However, Applicant does not necessarily concur with the interpretation of the dependent claims as set forth in the Office Action, nor does Applicant concur that the basis for the rejection of any of these claims is proper. Therefore, Applicant reserves the right to specifically address the patentability of these claims in the future.

12 General Comments on Dependent Claims

Since each of the dependent claims depends from a base claim that is believed to be in condition for allowance, Applicant believes that it is unnecessary at this time to argue the allowability of each of the dependent claims individually. However, Applicant does not

necessarily concur with the interpretation of the dependent claims as set forth in the Office Action, nor does Applicant concur that the basis for the rejection of any of the dependent claims is proper. Therefore, Applicant reserves the right to specifically address the patentability of the dependent claims in the future.

NVIDIA Ex1002 Page 668 of 1081

CONCLUSION

It is respectfully believed that all of the rejections set forth in the Office Action have been addressed. However, the absence of a reply to a specific rejection, objection, or comment set forth in the Office Action does not signify agreement with or concession of that rejection, objection, or comment. In addition, there may be reasons for patentability of any or all pending claims (or other claims) that have not been expressed. Furthermore, nothing in this paper should be construed as an intent to concede any issue with regard to any claim.

Applicant believes that the present application is now in condition for allowance. Favorable reconsideration of the application as amended is respectfully requested.

The Examiner is invited to contact the undersigned by telephone if it is felt that a telephone interview would advance the prosecution of the present application.

Dated: March 20, 2019

Respectfully submitted,

Smith Baluch LLP
376 Boylston St, Suite 401
Boston MA 02116

Phone: 617-947-7280

By /Christopher Max Colice/
Christopher Max Colice
Registration No. 65,634

IN THE CLAIMS:

Set forth below in ascending order, with status identifiers, is a complete listing of all claims currently under examination. Per 37 C.F.R. § 1.173, changes to any amended claims are indicated by [brackets] and/or underlining. New claims are indicated by underlining. This listing also reflects any cancellation and/or addition of claims.

1. (Twice Amended) A computer system, comprising:

a central processing unit to receive input data;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the [first] main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations yielding intermediate results; and

accelerator memory, operably coupled to the [graphic] at least one graphics processing unit, to store the results of the [plurality of sequential] sequence of computations; and

a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to transfer the at least a portion of the input data into the accelerator memory and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one [graphic] graphics processing unit.

2. The computer system of claim 1, wherein the central processing unit is configured to receive the input data in response to a user interaction.

3. The computer system of claim 1, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one graphics processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

4. The computer system of claim 1, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

5. The computer system of claim 1, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

6. The computer system of claim 1, wherein accelerator memory comprises:

a first memory bank to store parameters common to all of the computations in the sequence of computations; and

a second memory bank to store data specific to at least one computation in the sequence of computations.

7. The computer system of claim 1, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

8. The computer system of claim 1, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one graphics processing unit has begun to perform another sequence of computations.

9. The computer system of claim 8, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one graphics processing unit.

10. The computer system of claim 1, wherein the controller is configured to control execution of the sequence of computations by the at least one graphics processing unit.

11. (Amended) The computer system of claim 1, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central [processor] <u>processing</u> unit, to acquire the input data in real time.

12. A method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit

via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory, the method comprising:

(A) performing, by the GPU, the sequence of computations on a first portion of [the] input data so as to generate a first portion of [the] output data, intermediate computations in the sequence of computations yielding intermediate results;

(B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; and

(C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus.

13. The method of claim 12, further comprising:

storing the input data in the main memory in response to a user interaction.

14. The method of claim 12, further comprising:

receiving the input data at a first rate; and

wherein (A) comprises performing the sequence of computations at a second rate different than the first rate.

15. The method of claim 12, wherein (A) comprises:

generating an output representative of an output of at least one neuron in an artificial neural network.

16. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory without transferring any of the intermediate results of the plurality of sequential computations from the accelerator memory to the main memory so as to reduce data transfer via the bus.

17. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory after the GPU has begun to perform another sequence of computations.

18. The method of claim 17, wherein (C) further comprises:

initiating transfer of the second portion of the output data in parallel with performance of at least one computation in the other sequence of computations.

19. The method of claim 12, further comprising:

acquiring the input data in real time with at least one of a video camera, a microphone, or a cell recording electrode operably coupled to the CPU.

20. The method of claim 12, further comprising:

storing parameters common to all of the computations in the sequence of computations in a first memory bank in the accelerator memory; and

storing data specific to at least one computation in the sequence of computations in a second memory bank in the accelerator memory.

21. (New, Amended) An accelerator system comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data, the sequence of computations representing an artificial neural network, computations in the sequence of computations representing respective layers of neurons in an artificial neural network;

a memory, operably coupled to the GPU, to store the input data and the output data; and

a controller, operably coupled to the GPU and the memory, to transfer the input data into the memory and to transfer at least a portion of the output data from the memory during performance of the sequence of computations by the GPU.

22. (New) The accelerator system of claim 21, wherein the GPU is operably coupled to a central processing unit (CPU) of a computer system.

23. (New) The accelerator system of claim 21, wherein the GPU is mounted on an expansion card operably coupled to a motherboard of a computer system.

24. (New) The accelerator system of claim 21, further comprising:

a central processing unit (CPU), a system memory, and a non-volatile memory installed on a motherboard.

25. (New) The accelerator system of claim 24, wherein the GPU is operably coupled to the CPU via a system bus on the motherboard.

26. (New) The accelerator system of claim 25, wherein the system bus has a bandwidth of 3.2 GB/s to 15 GB/s.

27. (New) The accelerator system of claim 25, wherein the controller is configured to transfer the input data and the at least a portion of the output data via the system bus.

28. (New) The accelerator system of claim 25, further comprising:

a video card, operably coupled to the motherboard via a local bus, to transfer graphical information to a display.

29. (New) The accelerator system of claim 28, wherein the local bus has a lower bandwidth than the system bus.

30. (New) The accelerator system of claim 21, wherein the memory comprises:

a first memory partition to store parameters common to every computation in the sequence of computations.

31. (New) The accelerator system of claim 30, wherein the memory further comprises:

a second memory partition to store data specific to the computation in the sequence of computations.

32. (New) The accelerator system of claim 31, wherein the second memory partition comprises:

a first partition to hold external input data patterns;

a second partition to hold representations of internal variables;

a third partition to hold an input of the computation in the sequence of computations; and

a fourth partition to hold the output of the computation in the sequence of computations.

33. (New) The accelerator system of claim 31, wherein the second memory partition is configured to accumulate the output of the computation in the sequence of computations during execution, by the GPU, of computational elements of the computation in the sequence of computations.

34. (New) The accelerator system of claim 31, wherein the second memory partition is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to computational elements executed by the GPU during the computation in the sequence of computations.

35. (New) The accelerator system of claim 21, wherein the computation in the sequence of computations performed by the GPU comprises a plurality of computational elements.

36. (New) The accelerator system of claim 35, wherein the plurality of computational elements represents a layer of neurons in an artificial neural network.

37. (New) The accelerator system of claim 36, wherein all neurons in the layer of neurons are described by the same equation.

38. (New) The accelerator system of claim 21, wherein the memory is configured to store intermediate results of the sequence of computations and the controller is configured to transfer the output data from the memory to a main memory without transferring the intermediate results.

39. (New) The accelerator system of claim 21, wherein the controller is configured to initiate transfer of at least a portion of the input data and to transfer the at least a portion of the output data in parallel.

40. (New) The accelerator system of claim 21, wherein the input data is acquired by at least one of a video camera, a microphone, or a cell recording electrode.

41. (New) The accelerator system of claim 21, wherein the input data is loaded from disk.

42. (New, Amended) A method of performing a sequence of computations representing an artificial neural network, the method comprising:

receiving, at a central processing unit (CPU), first input data acquired from an external system;

performing, by a graphics processing unit (GPU), a first computation in the sequence of computations on the first input data to generate first output data, computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of a computation in the sequence of computations representing an output of a neuron in the artificial neural network;

storing, in a memory operably coupled to the GPU, the first input data and the first output data; and

transferring second input data into the memory during performance of the sequence of computations by the GPU.

43. (New, Amended) The method of claim 42, wherein transferring the second input data comprises transferring the second input data via a bus operably coupled to the central processing unit (CPU).

44. (New) The method of claim 42, further comprising:

transferring the first output data from the memory to another memory during a second computation in the sequence of computations.

45. (New) The method of claim 44, further comprising:

storing intermediate results of the sequence of computations in the memory, and

wherein transferring the first output data from the memory to the other memory occurs without transferring the intermediate results of the sequence of computations.

46. (New) The method of claim 44, wherein transferring the second input data and transferring the first output data occurs in parallel.

47. (New) The method of claim 42, further comprising:

storing, in a first memory partition of the memory, parameters common to all of the computations in the sequence of computations.

48. (New) The method of claim 47, further comprising:

storing, in a second memory partition of the memory, data specific to the first computation in the sequence of computations.

49. (New) The method of claim 48, further comprising:

storing, in the second memory partition, external input data patterns, representations of internal variables, an input of the computation in the sequence of computations, and the output of the computation in the sequence of computations.

50. (New) The method of claim 42, wherein storing the first output data comprises:

accumulating, in the memory, outputs of computational elements executed by the GPU in performing the first computation in the sequence of computations.

51. (New) The method of claim 42, further comprising:

storing, in the memory, an output of a previous computation in the sequence of computations; and

accessing, by the GPU, the output of the previous computation during performance of the computation in the sequence of computations.

52. (New) The method of claim 42, wherein performing the first computation comprises executing a plurality of computational elements representing a layer of neurons in an artificial neural network.

53. (New) The method of claim 52, wherein all neurons in the layer of neurons are described by the same equation.

54. (New) The method of claim 42, further comprising:

acquiring the input data with at least one of a video camera, a microphone, or a cell recording electrode.

55. (New, Amended) The method of claim 42, further comprising:

loading the input data from disk.

56. (New, Amended) An accelerator system comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network;

a memory, operably coupled to the GPU, to store parameters common to all of the computations in the sequence of computations and data specific to the computation in the sequence of computations; and

a controller, operably coupled to the GPU and the memory, to transfer at least a portion of the input data into the memory in parallel with performance of the sequence of computations

by the GPU, and to transfer at least a portion of the output data from the memory in parallel with performance of the sequence of computations by the GPU.

57. (New) The accelerator system of claim 56, wherein all neurons in the layer of neurons are described by the same equation.

58. (New) The accelerator system of claim 56, wherein the memory is configured to hold:
external input data patterns;
representations of internal variables;
an input to the computation in the sequence of computations; and
the output of the computation in the sequence of computations.

59. (New) The accelerator system of claim 56, wherein the memory is configured to accumulate the output of the computation in the sequence of computations during execution of the plurality of computational elements by the GPU.

60. (New) The accelerator system of claim 56, wherein the memory is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to the plurality of computational elements during the computation in the sequence of computations.

61. Cancel claim 61.

62. (New, Amended) A method comprising:
receiving, at a central processing unit (CPU), input data acquired from an external system;
performing, with a graphics processing unit (GPU), a sequence of computations on the input data to generate output data, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network;
storing, in a memory operably coupled to the GPU, parameters common to all of the computations in the sequence of computations;
storing, in the memory, data specific to the computation in the sequence of computations; and

transferring the input data into the memory and at least a portion of the output data from the memory during performance of the sequence of computations by the GPU.

63. (New) The method of claim 62, wherein all neurons in the layer of neurons are described by the same equation.

64. (New) The method of claim 62, further comprising:

storing, in the memory, external input data patterns;

storing, in the memory, representations of internal variables; and

storing, in the memory, an input of the computation in the sequence of computations.

65. (New) The method of claim 62, further comprising:

accumulating, in the memory, the output of the computation in the sequence of computations during execution of the plurality of computational elements by the GPU.

66. (New) The method of claim 62, further comprising:

storing, in the memory, an output of a previous computation in the sequence of computations; and

accessing, by the GPU, the output of the previous computation during the computation in the sequence of computations.

67. Cancel claim 67

68. (New, Amended) A method of executing computations representing an artificial neural network on a computer system comprising at least one central processing unit (CPU), a processing unit, a first memory partition, and a second memory partition, the method comprising:

executing, by the at least one CPU, a user interaction stream, the user interaction stream controlling transfer of inputs to the artificial neural network to the first memory partition and the second memory partition;

executing, by the processing unit, a computational stream, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network.

69. (New) The method of claim 68, wherein the user interaction stream controls data exchange between the user interaction stream and the computational stream outside of execution of the computations representing the artificial neural network.

70. (New) The method of claim 68, wherein executing the user interaction stream comprises:

controlling setting and editing of computational elements of the computations representing the artificial neural network.

71. (New) The method of claim 68, wherein executing the user interaction stream comprises:

controlling setting and editing of parameters of the computations representing the artificial neural network.

72. (New) The method of claim 68, wherein executing the user interaction stream comprises:

controlling setting and editing of parameters of the inputs to the artificial neural network.

73. (New) The method of claim 68, wherein executing the user interaction stream comprises:

specifying an output to be saved to disk and/or displayed on a screen.

74. (New) The method of claim 68, wherein executing the user interaction stream comprises:

parsing elements to be used in the computations representing the artificial neural network.

75. (New) The method of claim 74, wherein the processing unit comprises a graphics processing unit (CPU) and executing the user interaction stream further comprises:

converting the elements into GPU programs.

76. (New) The method of claim 75, wherein executing the user interaction stream comprises:

compiling the GPU programs.

77. (New) The method of claim 76, wherein executing the user interaction stream comprises:

transferring the GPU programs to the second memory partition.

78. (New) The method of claim 68, further comprising:

executing, by the at least one CPU, a data output stream, the data output stream controlling transfer of outputs of the computations representing the artificial neural network to disk.

79. (New, Amended) A system for executing computations representing an artificial neural network, the system comprising:

a first memory partition;

a second memory partition;

at least one central processing unit (CPU), operably coupled to the at least one memory, to execute a user interaction stream, the user interaction stream controlling transfer of inputs to the artificial neural network to the first memory partition and the second memory partition; and

a processing unit, operably coupled to the at least one memory and the at least one CPU, to execute a computational stream, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network.

80. (New) The system of claim 79, wherein the user interaction stream controls data exchange between the user interaction stream and the computational stream outside of execution of the computations representing the artificial neural network.

81. (New) The system of claim 79, wherein the user interaction stream controls setting and editing of computational elements of the computations representing the artificial neural network.

82. (New) The system of claim 79, wherein the user interaction stream controls setting and editing of parameters of the computations representing the artificial neural network.

83. (New) The system of claim 79, wherein the user interaction stream controls setting and editing of parameters of the inputs to the artificial neural network.

84. (New) The system of claim 79, wherein the user interaction stream specifies an output to be saved to disk and/or displayed on a screen.

85. (New) The system of claim 79, wherein the user interaction stream parses elements to be used in the computations representing the artificial neural network.

86. (New) The system of claim 85, wherein the processing unit comprises a graphics processing unit (GPU) and the user interaction stream converts the elements into GPU programs.

87. (New) The system of claim 86, wherein the user interaction stream compiles the GPU programs.

88. (New, Amended) The system of claim 87, wherein the user interaction stream transfers the GPU programs to at least one of the first memory partition or the second memory partition.

89. (New) The system of claim 68, wherein the at least one CPU is further configured to execute a data output stream, the data output stream controlling transfer of outputs of the computations representing the artificial neural network to disk.

90. (New, Amended) A method of performing a sequence of computations with a processing unit, the method comprising:

receiving input data;

performing, by the processing unit, the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network;

in parallel with performing the sequence of computations by the processing unit, transferring a second portion of the input data from a first memory partition to a second memory partition; and

in parallel with performing the sequence of computations by the processing unit, transferring a second portion of the output data from the second memory partition to the first memory partition.

91. (New) The method of claim 90, wherein receiving the input data comprises loading the input data from disk.

92. (New) The method of claim 90, wherein receiving the input data comprises receiving the input data from at least one of a video camera, a microphone, or a cell recording electrode.

93. (New) The method of claim 90, wherein receiving the input data occurs at a first rate and performing the sequence of computations occurs at a second rate different than the first rate.

94. (New) The method of claim 90, wherein transferring the second portion of the output data comprises transferring the second portion of the output data from the second memory partition to the first memory partition without transferring any of the intermediate results from the second memory partition to the first memory partition.

95. (New, Amended) A system for performing a sequence of computations, the system comprising:

a first memory partition;

a second memory partition operably coupled to the first memory partition; and

a processing unit, operably coupled to the first memory partition and the second memory partition, to perform the sequence of computations on a first portion of input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network,

wherein the first memory partition is configured to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit, and

wherein the second memory partition is configured to transfer a second portion of the output data to the first memory partition in parallel with performance the sequence of computations by the processing unit.

96. (New) The system of claim 95, wherein the first memory partition and the second memory partition are logical partitions.

97. (New) The system of claim 95, wherein the processing unit is comprises a graphics processing unit (GPU).

98. (New) The system of claim 95, wherein the processing unit is configured to receive the input data at a first rate and to perform the sequence of computations at a second rate is different than the first rate.

99. (New) The system of claim 95, wherein the second memory partition is configured to transfer the second portion of the output data to the first memory partition without transferring any of the intermediate results to the first memory partition.

100. (New) The system of claim 95, further comprising

at least one of a video camera, a microphone, or a cell recording electrode to generate the input data.

101. (New, Amended) A method for performing a numerical simulation representing an artificial neural network over a plurality of computational cycles including at least a first computational cycle representing a first layer of neurons in the artificial neural network and a second computational cycle representing a second layer of neurons in the artificial neural network, the method comprising:

storing, in a first partition of a memory, first output data during the second computational cycle, the first output data generated during the first computational cycle and being referenced by a first pointer, the first output data representing at least one output of the first layer of neurons;

performing, by a processing unit operably coupled to the memory, at least one calculation on the first output data during the second computational cycle, the at least one calculation yielding second output data, the second output data representing at least one output of the second layer of neurons;

storing, in a second partition of the memory, the second output data, the second output data being referenced by a second pointer; and

swapping the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles.

102. (New) The method of claim 101, further comprising:

transferring input data into the memory from another memory coupled to a central processing unit (CPU) before the first computational cycle.

103. (New) The method of claim 101, further comprising:

providing, to a central processing unit (CPU), first output data generated during the first computational cycle after a conclusion of the first computational cycle.

104. (New) The method of claim 101, further comprising:

transferring the first output data from the memory to another memory during the second computational cycle.

105. (New) The method of claim 101, further comprising:

displaying, with a video system, an indication of the numerical simulation.

106. (New) The method of claim 101, further comprising:

dictating, to the processing unit, an order of execution of instructions.

107. (New) The method of claim 101, wherein performing the at least one calculation comprises receiving, by the processing unit, at least the portion of the input data and a series of instructions and executing, by the processing unit, the instructions on the at least the portion of the input data.

108. (New, Amended) A method of performing a numerical simulation representing an artificial neural network, the method comprising:

performing, by a processing unit, at least one computation during a first computational cycle of the numerical simulation, the first computational cycle representing a first layer of neurons in the artificial neural network;

storing, by a memory operably coupled to the processing unit, data used to perform the at least one computation;

controlling performance of the at least one computation by the processing unit;

controlling transfer of the output data from the processing unit to the memory during the first computational cycle; and

controlling transfer of input data to the memory for use by the processing unit during a second computational cycle of the numerical simulation, the second computational cycle representing a second layer of neurons in the artificial neural network.

109. (New) The method of claim 108, further comprising:

dictating an order of execution of instructions to the processing unit.

110. (New) The method of claim 108, further comprising:

transferring output data from the memory to a main memory while the processing unit performs the at least one computation.

111. (New) The method of claim 108, further comprising:

perform successive computational cycles of the numerical simulation by feeding output data generated by the processing unit from a previous computational cycle and input data for a next computational cycle into the processing unit.

112. (New, Amended) A system for executing an artificial neural network, the system comprising:

a central processing unit (CPU) to provide first input data;

a memory, operably coupled to the CPU, to store the first input data in a first partition, referenced by a first pointer, before computing a first layer of neurons of the artificial neural network;

a processing unit, operably coupled to the memory, to perform, during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data, the first output data representing an output of at least one neuron in the first layer of neurons; and

a controller, operably coupled to the processing unit and the memory, to store the first output data in a second partition of the memory, the second partition referenced by a second pointer, and to swap the first pointer with the second pointer at the end of the computation of the first layer of neurons, such that the first output data becomes an input for a second layer of neurons of the artificial neural network.

113. (New) The system of claim 112, wherein the processing unit comprises a graphics processing unit.

114. (New) The system of claim 112, wherein the controller is configured to send instructions for performing the at least one calculation to the processing unit.

115. (New, Amended) The system of claim 112, wherein the controller is further configured to transfer the first output data to another memory during computation of the second layer of neurons.

116. (New, Amended) The system of claim 112, wherein the memory further comprises:

a third partition to store internal variables; and

a fourth partition to store data used as input at a particular layer of neurons of the artificial neural network.

117. (New, Amended) A computer system, comprising:

a central processing unit to receive input data acquired from an external system;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations yielding intermediate results; and

accelerator memory, operably coupled to the at least one processing unit, to store the results of the sequence of computations; and

a controller, operably coupled to the at least one processing unit and the accelerator memory, to transfer the at least a portion of the input data into the accelerator memory and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one processing unit.

118. (New) The computer system of claim 117, wherein the central processing unit is configured to receive the input data in response to a user interaction.

119. (New) The computer system of claim 117, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

120. (New) The computer system of claim 117, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

121. (New) The computer system of claim 117, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

122. (New) The computer system of claim 117, wherein accelerator memory comprises:

a first memory partition to store parameters common to all of the computations in the sequence of computations; and

a second memory partition to store data specific to at least one computation in the sequence of computations.

123. (New) The computer system of claim 117, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

124. (New) The computer system of claim 117, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one processing unit has begun to perform another sequence of computations.

125. (New) The computer system of claim 124, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one processing unit.

126. (New) The computer system of claim 117, wherein the controller is configured to control execution of the sequence of computations by the at least one processing unit.

127. (New, Amended) The computer system of claim 117, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time.

128. (New, Amended) A method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a processing unit and an accelerator memory, the method comprising:

(A) performing, by the processing unit, the sequence of computations on a first portion of input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results;

(B) in parallel with performing the sequence of computations by the processing unit in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; and

(C) in parallel with performing the sequence of computations by the processing unit in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus.

129. (New) The method of claim 128, further comprising:

storing the input data in the main memory in response to a user interaction.

130. (New) The method of claim 128, further comprising:

receiving the input data at a first rate; and

wherein (A) comprises performing the sequence of computations at a second rate different than the first rate.

131. (New) The method of claim 128, wherein (A) comprises:

generating an output representative of an output of at least one neuron in an artificial neural network.

132. (New) The method of claim 128, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory without transferring any of the intermediate results of the plurality of sequential computations from the accelerator memory to the main memory so as to reduce data transfer via the bus.

133. (New) The method of claim 128, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory after the processing unit has begun to perform another sequence of computations.

134. (New) The method of claim 133, wherein (C) further comprises:

initiating transfer of the second portion of the output data in parallel with performance of at least one computation in the other sequence of computations.

135. (New) The method of claim 128, further comprising:

acquiring the input data in real time with at least one of a video camera, a microphone, or a cell recording electrode operably coupled to the processing unit.

136. (New) The method of claim 128, further comprising:

storing parameters common to all of the computations in the sequence of computations in a first memory partition in the accelerator memory; and

storing data specific to at least one computation in the sequence of computations in a second memory partition in the accelerator memory.

137. (New) A system for executing an artificial neural network, the system comprising:

a first processing unit to receive input data in real time from a sensor;

a memory, operably coupled to the first processing unit, to store a first portion of the input data;

a second processing unit, operably coupled to the memory, to generate first output data from the first portion of the input data, the first output data representing an output of at least one neuron in a first layer of neurons in the artificial neural network;

a controller, operably coupled to the second processing unit and the memory, to store a second portion of the input data in the memory while the second processing unit generates the first output data.

138. (New) The system of claim 137, wherein the sensor is a microphone.

139. (New) The system of claim 137, wherein the sensor is a video camera.

140. (New) The system of claim 137, wherein the second processing unit is further configured to generate second output data from the second portion of the input data, the second output data representing an output of at least one neuron in a second layer of neurons in the artificial neural network.

141. (New) The system of claim 140, further comprising:

a random access memory, operably coupled to the first processing unit, to receive the first output data from the memory while the second processing unit generates the second output data.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 15/808,201	Filing Date 11/09/2017	☐ To be Mailed

ENTITY: ☐ LARGE ☑ SMALL ☐ MICRO

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A		N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A		N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*		x $40 =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*		x $210 =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).				
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))					
* If the difference in column 1 is less than zero, enter "0" in column 2.				TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT	03/21/2019	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 136	Minus	** 136	= 0		x $50 =	0
	Independent (37 CFR 1.16(h))	* 15	Minus	*** 15	= 0		x $230 =	0
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=		x $0 =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=		x $0 =	
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.	LIE
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".	/PATRICIA F LEWIS/
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".	
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.	

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 15/808,201	Filing Date 11/09/2017	☐ To be Mailed

ENTITY: ☐ LARGE ☑ SMALL ☐ MICRO

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A		N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A		N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*		x $40 =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*		x $210 =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).				
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))					
* If the difference in column 1 is less than zero, enter "0" in column 2.				TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT	03/21/2019	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 139	Minus	** 136	= 3		x $50 =	150
	Independent (37 CFR 1.16(h))	* 16	Minus	*** 15	= 1		x $230 =	230
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	380

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=		x $0 =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=		x $0 =	
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.	LIE
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".	/PATRICIA F LEWIS/
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".	
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.	

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: REIS.DECL
Document Description: Reissue Declaration Filed In Accordance With MPEP 1414

PTO/AIA/05 (06-12)
Approved for use through 08/31/2013. OMB 0651-0033
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

REISSUE APPLICATION DECLARATION BY THE INVENTOR	Docket Number (Optional) NRLA-001US03

I hereby declare that:
Each inventor's residence and mailing address are stated below next to their name.
I believe I am the original inventor or an original joint inventor of the subject matter which is described and claimed in patent number 9,189,828, granted November 17, 2015 and for which a reissue patent is sought on the invention titled Graphic Processor Based Accelerator System and Method,

the specification of which

☐ is attached hereto.

☑ was filed on 11/9/2017 as reissue application number 15/808,201 .

The above-identified application was made or authorized to be made by me.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

I believe the original patent to be wholly or partly inoperative or invalid, for the reasons described below. (Check all boxes that apply.)

☐ by reason of a defective specification or drawing.

☑ by reason of the patentee claiming more or less than he had the right to claim in the patent.

☑ by reason of other errors.

At least one error upon which reissue is based is described below. If the reissue is a broadening reissue, a claim that the application seeks to broaden must be identified:

Reissue is requested because of antecedent basis errors in issued claim 1 of US 9,189,828. Specifically, claim 1 lacks antecedent basis for "the plurality of computations," which is being amended to recite "the sequence of computations" to correct this error. Reissue is also requested because Applicant erred by claiming less than it had a right to claim in US 9,189,828. Claim 21, for example, corrects this error because it does not recite "a central processing unit" as in claim 1. And claim 79, for example, corrects Applicant's error by claiming the "user interaction stream" that appears in the embodiment of FIG. 3 of US 9,189,828 and has not been claimed to date.

This collection of information is required by 37 CFR 1.175. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

(REISSUE APPLICATION DECLARATION BY THE INVENTOR, page 2)

Docket Number (Optional): NRLA-001US03

Note: To appoint a power of attorney, use form PTO/AIA/81.

Correspondence Address: Direct all communications about the application to:

☑ The address associated with Customer Number: 149890

OR

☐ Firm or Individual Name	
Address	
City	State / Zip
Country	
Telephone	Email

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

Legal name of sole or first inventor (*E.g.*, Given Name (first and middle (if any) and Family Name or Surname)

Fabrizio Santini

Inventor's Signature	Date (Optional)
Residence: City / State / Country	
Mailing Address	
City / State / Zip / Country	

☐ Additional joint inventors are named on the _____ supplemental sheet(s) PTO/AIA/10 attached hereto.

NVIDIA Ex1002 Page 694 of 1081

Electronic Acknowledgement Receipt

EFS ID:	35489124
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice/Lindsey Lovato
Filer Authorized By:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03
Receipt Date:	21-MAR-2019
Filing Date:	09-NOV-2017
Time Stamp:	14:20:58
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Oath or Declaration filed	NRLA-001US03_Fabrizio_DEC.pdf	663266 ab3ab15634015aec57ec445de780d6a970daaa87	no	2

Warnings:

Information:	
Total Files Size (in bytes):	663266

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 05/13/2019

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

NOTIFICATION DATE	DELIVERY MODE
05/13/2019	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docketing@smithbaluch.com
smith@smithbaluch.com

Office Action Summary	Application No. 15/808,201	Applicant(s) Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	AIA (FITF) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☑ Responsive to communication(s) filed on 03/20/2019.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☑ This action is **FINAL.** 2b) ☐ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5) ☑ Claim(s) 1-60,62-66 and 68-141 is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6) ☐ Claim(s) _____ is/are allowed.

7) ☑ Claim(s) 1-60,62-66 and 68-141 is/are rejected.

8) ☐ Claim(s) _____ is/are objected to.

9) ☐ Claim(s) _____ are subject to restriction and/or election requirement

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to **PPHfeedback@uspto.gov.**

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)

2) ☑ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date 03/20/2019.

3) ☑ Interview Summary (PTO-413) Paper No(s)/Mail Date 20190418.

4) ☐ Other: _____.

The present application is being examined under the pre-AIA first to invent provisions.

DETAILED ACTION

For reissue applications filed before September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the law and rules in effect on September 15, 2012. Where specifically designated, these are "pre-AIA" provisions.

For reissue applications filed on or after September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the current provisions.

Claims 1-60, 62-66, and 68-141 are currently pending in the application. Claims 1-20 are original claims to patent US 9,189,828 B2 and claims 21-60, 62-66, and 68-141 are newly added claims.

Reissue

Applicant is reminded of the continuing obligation under 37 CFR 1.178(b), to timely apprise the Office of any prior or concurrent proceeding in which Patent No. 9,189,828 is or was involved. These proceedings would include interferences, reissues, reexaminations, and litigation.

Applicant is further reminded of the continuing obligation under 37 CFR 1.56, to timely apprise the Office of any information which is material to patentability of the claims under consideration in this reissue application.

These obligations rest with each individual associated with the filing and prosecution of this application for reissue. See also MPEP §§ 1404, 1442.01 and 1442.04.

Applicant is notified that any subsequent amendment to the specification and/or claims must comply with 37 CFR 1.173(b).

Information Disclosure Statement

The information disclosure statement (IDS) submitted on 03/20/2019 have been considered in accordance with MPEP 609, including timing and content requirements (see also 37 C.F.R. 1.97 and 1.98). Documents not meeting a particular criteria are lined through and not considered. NOTE: Item 1 of the Non-Patent Literature is titled "Brain-inspired computing". No such document could be found. However, one submitted document does not appear on the IDS and is titled "From Spikes to Inter-Areal Synchrony: How attentive matching and resonance control learning and information processing by laminar thalamocortical circuits".

Claim Objections

Claim 12 is objected to because of the following informalities: the amendment (03/20/2019) does not contain a parenthetical expression as required by 1.173(b)(2) or a status as required by CFR 1.173(c). Appropriate correction is required.

Claims 21-41 are objected to because of the following informalities: the amendment (03/20/2019) does not provide support for the changes as required by CFR 1.173(c). The support provided (on page 25 of the response) does not show computations represent layers. Appropriate correction is required.

Claims 42-54 are objected to because of the following informalities: the amendment (03/20/2019) does not provide support for the changes as required by CFR 1.173(c). No support is provided in any manner. Appropriate correction is required.

Claims 42-54, 62-66, and 117-127 are objected to because of the following informalities: the amendment (03/20/2019) does not provide support for the changes as required by CFR 1.173(c), specifically the claim language "external system". Appropriate correction is required.

Claims 101-116 and 137-141 are objected to because of the following informalities: the amendment (03/20/2019) does not provide support for the changes as required by CFR 1.173(c),

specifically the claim language relating computational cycles to first and second layers of neurons.

Appropriate correction is required.

Claim Rejections - 35 USC § 112

The following is a quotation of 35 U.S.C. 112(b):
(b) CONCLUSION.—The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the inventor or a joint inventor regards as the invention.

The following is a quotation of 35 U.S.C. 112 (pre-AIA), second paragraph:
The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claims 62-66 are rejected under 35 U.S.C. 112(b) or 35 U.S.C. 112 (pre-AIA), second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which the inventor or a joint inventor, or for pre-AIA the applicant regards as the invention. In particular, independent claim 62, upon which claims 63-66 depend, now states, "transferring the input data into the memory … during performance of the sequence of computations by the GPU". This limitation is treated as stating, "transferring at least a portion of the input data …". This interpretation is necessary in order for at least a portion of the input data to have already been transferred so that in can be performed upon by the GPU.

Claim Analysis – 35 USC § 251

Withdrawn are the rejections of claims 21-116 under 35 U.S.C. 251 as being an improper recapture of broadened claimed subject matter surrendered in the application for the patent upon which the present reissue is based.

It is noted that the following is the three step test for determining recapture in reissue applications (see: MPEP 1412.02(II)): "(1) first, we determine whether, and in what respect, the reissue

claims are broader in scope than the original patent claims; (2) next, we determine whether the broader aspects of the reissue claims relate to subject matter surrendered in the original prosecution; and (3) finally, we determine whether the reissue claims were materially narrowed in other respects, so that the claims may not have been enlarged, and hence avoid the recapture rule."

Claims directed toward separate inventions/embodiments/species that were disclosed but never covered by the original prosecution claims may be to overlooked aspects. Claims to overlooked aspects are not subject to recapture because the claims are, by definition, unrelated to subject matter that was surrendered during the prosecution of the original application (see MPEP 1412.01 II. and MPEP 1402.02).

Upon further consideration, including in view of Patent Owner's response of 03/20/2019, the following concise analysis is provided and supersedes the analysis of the preceding office action of 11/20/2018.

Analysis 1: The record of the prior proceedings 11/860,254 (US 8,648,867 B2) and 14/147,015 (US 9,189,828 B2) indicate:

a. In application 14/147,015, the Amendment of 03/10/2015 added to claim 1 (issued as claim 1 in patent 9,189,828) the following limitations in order to overcome the cited prior art (additions shown by underlining): a central processing unit to receive input data; main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit; an accelerator, operably coupled to the central processing unit and the first memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising: at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, intermediate computations in

the sequence of computations yielding intermediate results; and accelerator memory, operably coupled to the graphic processing unit, to store the results of the plurality of sequential computations; and a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to transfer the at least a portion of the input data into the accelerator memory and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one graphic processing unit. The accompanying remarks of 03/10/2015, on page 8, further indicate the following limitations overcome the cited prior art: "a controller … to transfer at least a portion of the output data from the accelerator memory to [a] main memory during performance of [a] sequence of computations by [a] graphic processing unit".

b. In application 14/147,015, the Amendment of 03/10/2015 added claim 12 (issued as claim 12 in patent 9,189,828) with the following limitations in order to overcome the cited prior art (additions shown by underlining): A method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory, the method comprising: (A) performing, by the GPU, the sequence of computations on a first portion of the input data so as to generate a first portion of the output data, intermediate computations in the sequence of computations yielding intermediate results; (B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; and (C) in parallel with performing the sequence of computations by the GPU in (A),

transferring a second portion of the output data from the accelerator memory to the main memory via the bus.

c. In application 11/860,254, the Amendment of 07/31/2013 added to claim 1 (issued as claim 1 in patent 8,648,867) the following limitations in order to overcome the cited prior art (additions shown by underlining): A computer system for performing a numerical simulation over a plurality of computational cycles including at least a first computational cycle and a second computational cycle, the computer system comprising: a central processing unit; a main memory, operably coupled to the central processing unit, to store input data to be accessed by the central processing unit in performing the numerical simulation; a video system, operably coupled to the central processing unit, to drive a video monitor to display an indication of the numerical simulation in response to the computer system performing the numerical simulation; an accelerator, operably coupled to the central processing unit, to receive at least a portion of the input data from the central processing unit and to provide first output data generated during the first computational cycle to the central processing unit after a conclusion of the first computational cycle, the accelerator comprising: at least one graphics processing unit to generate second output data, during the second computational cycle, by performing at least one calculation on the first output data; and an accelerator memory, operably coupled to the at least one graphic processing unit, the accelerator memory comprising: a first partition, referenced by a first pointer, to store the first output data during the second computational cycle; and a second partition, referenced by a second pointer, to store the second output data generated during the second computational cycle; and an accelerator controller, operably coupled to the accelerator memory and the central processing unit, to transfer the at least the

portion of the input data into the accelerator memory before the first computational cycle, to transfer the first output data from the accelerator memory to the main memory during the second computational cycle, to direct the second output data into the second partition during the second computational cycle, and to swap the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles. The accompanying remarks of 07/31/2013, on pages 11-13, further indicate the following limitations overcome the cited prior art: "accelerator controller, operably coupled to the accelerator memory and the central processing unit, to transfer the at least the portion of the input data into the accelerator memory before the first computational cycle, to transfer the first output data from the accelerator memory to the main memory during the second computational cycle, to direct the second output data into the second partition at the end of the second computational cycle, and to swap the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles". The remarks argue claims 1 and 17 are similar in this regard.

d. In application 11/860,254, the Amendment of 07/31/2013 added to claim 10 (issued as claim 9 in patent 8,648,867) the following limitations in order to overcome the cited prior art (additions shown by underlining): An accelerator system for a computer system performing a numerical simulation, the accelerator system comprising: at least one graphics processing unit to generate output data by performing at least one computation during a first computational cycle of the numerical simulation; an accelerator memory, operably coupled to the at least one graphics processing unit, to

store data used to perform the at least one computation; and an accelerator controller, operably coupled to the accelerator memory and the at least one graphics processing unit, to execute: (i) a computational stream controlling performance of the at least one computation by the at least one graphics processing unit; (ii) an output stream controlling transfer of the output data from the at least one graphics processing unit to the accelerator memory during the first computational cycle; and (iii) an input stream controlling transfer of input data to the accelerator memory for use by the at least one graphics processing unit during a second computational cycle of the numerical simulation. The accompanying remarks of 07/31/2013, on page 14, further indicate the following limitations overcome the cited prior art: an accelerator controller that executes a computational stream to control performance of a first computation, an output stream to transfer data at the end of the first computational cycle, and an input stream to control the transfer of input data into the accelerator memory.

e. In application 11/860,254, the Amendment of 07/31/2013 added to claim 17 (issued as claim 16 in patent 8,648,867) the following limitations in order to overcome the cited prior art (additions shown by underlining): A method for performing a numerical simulation on input data in a computer system including a central processing unit and an accelerator, the method comprising: receiving, by an accelerator, first input data from the central processing unit; transferring, by an accelerator controller, the first input data into a first partition, referenced by first pointer, of an accelerator memory before a first computational cycle of the numerical simulation; performing, by at least one graphics processing unit during the first computational cycle, at least one calculation on the first portion of the input data as to generate first output data; storing, by the accelerator controller, the first output data into a second partition,

referenced by a second pointer, of the accelerator memory; and swapping the first pointer with the second pointer at the end of the first computational cycle, such that the first output data becomes an input for a second computational cycle of the numerical simulation.

Analysis 2: Claims directed toward separate inventions/embodiments/species that were disclosed but never covered by the original prosecution claims may be to overlooked aspects. Claims to overlooked aspects are not subject to recapture because the claims are, by definition, unrelated to subject matter that was surrendered during the prosecution of the original application (see MPEP 1412.01 II. and MPEP 1402.02). The following independent claims are subject to recapture analysis as follows:

a. Claim 21: no recapture; the instant claim is directed toward an accelerator system with the closest comparison being with claim 9 of USPN 8,648,867 (see above Analysis 1d), however that claim is directed toward numerical simulation making use of various streams; though "numerical simulation" is broad, the instant claim 21 is further narrowed by "the sequence of computations representing an artificial neural network, computations in the sequence of computations representing respective layers of neurons in an artificial neural network"; in summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

b. Claim 42: no recapture; the instant claim is directed toward "a method of performing a sequence of computations representing an artificial neural network", but does not state the main memory found in claims 1 and 12 of USPN 9,189,828; the instant claim is further narrowed to "computations in the sequence of computations representing respective layers of neurons in

the artificial neural network"; in summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

c. Claim 56: no recapture; the instant claim is directed toward an accelerator system, does not have a main memory, and is not directed toward various "streams"; the instant claim is further narrowed to "computations in the sequence of computational elements representing a layer of neurons in an artificial neural network, an output of the computation in the sequence of computations representing an output of a neuron in the artificial network"; in summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

d. Claim 62: no recapture for similar reasons as claim 42;

e. Claim 68: no recapture; the instant claim is directed toward "a method of executing computations representing an artificial neural network on a computer system comprising at least one central processing unit" with the closest comparison being with claim 16 of USPN 8,648,867 (see above Analysis 1e), however that claim is directed toward numerical simulation; the instant claim 68 is narrowed by "an artificial neural network"; in summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

f. Claim 79: no recapture; the instant claim is directed toward "a system for executing computations representing an artificial neural network" with the closest comparison being with claim 9 of USPN 8,648,867 (see above Analysis 1d), however that claim is directed toward numerical simulation; the instant claim 79 is narrowed by "an artificial neural network"; in

summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

g. Claim 90: no recapture; the instant claim is directed toward "a method of performing a sequence of computations", but does not state the CPU or main memory found in claims 1 and 12 of USPN 9,189,828; the instant claim is further materially narrowed by "the first portion of the output data representing an output of an artificial neural network"; in summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

h. Claim 95: no recapture for similar reasons as to claim 90;

i. Claim 101: no recapture; the instant claim is directed toward a method for performing a numerical simulation representing an artificial neural network with the closest comparison being with claims 1 and 16 of USPN 8,648,867 (see above Analysis 1c and 1e); the instant claim 101 is further narrowed by "representing a first layer of neurons in the artificial neural network" and "representing a second layer of neurons in the artificial neural network"; in summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

j. Claim 108: no recapture for similar reasons as to claim 101;

k. Claim 112: no recapture; the instant claim is directed to a "system for executing an artificial neural network" further limited such that first output data represents an output in a first layer of neurons and then becomes an input for a second layer of neurons; the claimed "processing unit" is in reference to an accelerator/gpu which previous prosecution claims did not add to

overcome prior art (see above Analysis 1a, 1c and 1d); in summary the claim is now constituted to different aspects as compared to the prosecution claims, and as such the removed limitations are not to surrendered subject matter, and further the claim is materially narrowed;

l. Claims 117 and 128: no recapture as indicated in the previous office action of 11/20/2018;

m. Claim 137: no recapture; the instant claim is directed to a "system for executing an artificial neural network" further limited such that first output data represents an output in a first layer of neurons; this further limitation negates the need for the previous prosecution added limitations "intermediate computations in the sequence of computations yielding intermediate results" and "to transfer the at least a portion of".

Claim Rejections - 35 USC § 102 and § 103

In the event the determination of the status of the application as subject to AIA 35 U.S.C. 102 and 103 (or as subject to pre-AIA 35 U.S.C. 102 and 103) is incorrect, any correction of the statutory basis for the rejection will not be considered a new ground of rejection if the prior art relied upon, and the rationale supporting the rejection, would be the same under either status.

The following is a quotation of the appropriate paragraphs of pre-AIA 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (a) the invention was known or used by others in this country, or patented or described in a printed publication in this or a foreign country, before the invention thereof by the applicant for a patent.
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States

> only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

The following is a quotation of pre-AIA 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claim(s) 1-7, 10-16, 19-40, 42-43, 47-54, 56-60, 62-66, 68-72, 74-77, 79-83, 85-88, 90, 92-95, 97-100, and 117-123, 126-132, 135-138, and 140-141 is/are rejected under pre-AIA 35 U.S.C. 102(e) as being anticipated by US 7,219,085 B2 *(herein **Buck**)*.

Claim 1

Buck discloses a computer system *(Buck: figure 1)*, comprising:

a central processing unit to receive input data *(Buck: figure 1, elements 120 and 160; figure 2, element 120; column 9, lines 33-36)*;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit *(Buck: figure 1, elements 121, 130, 141, 152, 156)*;

an accelerator, operably coupled to the central processing unit and the main memory via the bus *(Buck: figure 1, element 190)*, to receive at least a portion of the input data from the main memory *(Buck: column 21, lines 54-64)*, the accelerator comprising:

at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, intermediate computations in the sequence of

computations yielding intermediate results *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)*; and

accelerator memory, operably coupled to the at least one graphics processing unit, to store the results of the sequence of computations *(Buck: figures 2 and 5, element 230)*; and

a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to transfer the at least a portion of the input data into the accelerator memory and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one graphics processing unit *(Buck: column 11, line 12 to column 12, line 26)*.

Claim 2

Buck discloses the computer system of claim 1, wherein the central processing unit is configured to receive the input data in response to a user interaction *(Buck: column 7, lines 51-55; column 23, lines 13-15)*.

Claim 3

Buck discloses the computer system of claim 1, wherein: the central processing unit is configured to receive the input data at a first rate; and the at least one graphics processing unit is configured to perform the sequence of computations at a second rate different than the first rate *(Buck: column 2, lines 55-67; column 9, lines 10-14; column 14, lines 40-44; CPU and GPU function at different rates)*.

Claim 4

Buck discloses the computer system of claim 1, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory *(Buck: column 7, lines 39-51; figure 5, elements 580 and 585)*.

Claim 5

Buck discloses the computer system of claim 1, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*.

Claim 6

Buck discloses the computer system of claim 1, wherein accelerator memory comprises: a first memory bank to store parameters common to all of the computations in the sequence of computations; and a second memory bank to store data specific to at least one computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Claim 7

Buck discloses the computer system of claim 1, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W; figure 5,*

temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W).

Claim 10

Buck discloses the computer system of claim 1, wherein the controller is configured to control execution of the sequence of computations by the at least one graphics processing unit *(Buck: column 12, lines 16-26)*.

Claim 11

Buck discloses the computer system of claim 1, further comprising: at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 23, lines 26-30)*.

Claim 12

Buck discloses a method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory *(Buck: figure 1, elements 120 and 160; figure 2, element 120; column 9, lines 33-36; figure 1, elements 121, 130, 141, 152, 156; figure 1, element 190; figure 2, elements 220 and 230)*, the method comprising:

(A) performing, by the GPU, the sequence of computations on a first portion of input data so as to generate a first portion of output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines*

16-26), intermediate computations in the sequence of computations yielding intermediate results *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)*;

(B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus *(Buck: column 11, line 12 to column 12, line 26)*; and

(C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus *(Buck: column 11, line 12 to column 12, line 26)*.

Claims 13-16 and 19-20

The limitations of claims 13-16 and 19-20 correspond to the limitations of claims 2-7 and 10-11 and as such are rejected in a corresponding manner.

Claim 21

Buck discloses an accelerator system *(Buck: figure 2, element 190)* comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, the sequence of computations representing an artificial neural network, computations in the sequence of computations representing layers of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*;

a memory, operably coupled to the GPU, to store the input data and the output data *(Buck: figure 5, element 230)*; and

a controller, operably coupled to the GPU and the memory, to transfer the input data into the memory and to transfer at least a portion of the output data from the memory during performance of

the sequence of computations by the GPU *(Buck: figure 5, inputs and outputs; column 11, line 59 to column 12, line 6)*.

Claim 22

Buck discloses the accelerator system of claim 21, wherein the GPU is operably coupled to a central processing unit (CPU) of a computer system *(Buck: figure 2)*.

Claim 23

Buck discloses the accelerator system of claim 21, wherein the GPU is mounted on an expansion card operably coupled to a motherboard of a computer system *(Buck: figure 2, element 210; column 9, lines 15-22, AGP and PCI Express busses are for expansion cards coupled to a motherboard)*.

Claim 24

Buck discloses the accelerator system of claim 21, further comprising: a central processing unit (CPU) *(Buck: figure 1, element 120)*, a system memory *(Buck: figure 1, element 130)*, and a non-volatile memory installed on a motherboard *(Buck: figure 1, element 133; column 6, lines 27-40, at least ISA bus indicates motherboard)*.

Claim 25

Buck discloses the accelerator system of claim 24, wherein the GPU is operably coupled to the CPU via a system bus on the motherboard *(Buck: figure 1, element 190 containing the GPU; figure 2, elements 190 and 210; column 9, lines 15-22, AGP and PCI Express busses are for expansion cards coupled to a motherboard)*.

Claim 26

Buck discloses the accelerator system of claim 25, wherein the system bus has a bandwidth of 3.2 GB/s to 15 GB/s *(Buck: column 9, lines 15-22, PCI Express bus has a bandwidth higher than 3.2GB/s)*.

Claim 27

Buck discloses the accelerator system of claim 25, wherein the controller is configured to transfer the input data and the at least a portion of the output data via the system bus *(Buck: figure 2, element 210)*.

Claim 28

Buck discloses the accelerator system of claim 25, further comprising: a video card, operably coupled to the motherboard via a local bus, to transfer graphical information to a display *(Buck: figure 1, elements 190, 121, and 191; claim language "local bus" is reference language "system bus")*.

Claim 29

Buck discloses the accelerator system of claim 28, wherein the local bus has a lower bandwidth than the system bus *(Buck: figure 1, element 121; column 6, lines 27-40, at least ISA bus; claim language "local bus" is reference language "system bus"; figure 2, element 210; column 9, lines 15-22; ISA bus has a lower bandwidth than a PCI Express bus)*.

Claim 30

Buck discloses the accelerator system of claim 21, wherein the memory comprises: a first memory partition to store parameters common to every computation in the sequence of computations *(Buck:*

figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4).

Claim 31

Buck discloses the accelerator system of claim 30, wherein the memory further comprises: a second memory partition to store data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Claim 32

Buck discloses the accelerator system of claim 31, wherein the second memory partition comprises: a first partition to hold external input data patterns *(Buck: figure 5, element 530; column 10, lines 61-62)*; a second partition to hold representations of internal variables *(Buck: column 15, lines 49-60)*; a third partition to hold an input of the computation in the sequence of computations *(Buck: column 10, lines 43-44)*; and a fourth partition to hold the output of the computation in the sequence of computations *(Buck: figure 5, element 510)*.

Claim 33

Buck discloses the accelerator system of claim 31, wherein the second memory partition is configured to accumulate the output of the computation in the sequence of computations during execution, by the GPU, of computational elements of the computation in the sequence of computations *(Buck: column 11, lines 50-52)*.

Claim 34

Buck discloses the accelerator system of claim 31, wherein the second memory partition is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to computational elements executed by the GPU during the computation in the sequence of computations *(Buck: figure 5, at least element 510; column 11, line 12 to column 12, line 15)*.

Claim 35

Buck discloses the accelerator system of claim 21, wherein the computation in the sequence of computations performed by the GPU comprises a plurality of computational elements *(Buck: column 21, lines 61-64)*.

Claim 36

Buck discloses the accelerator system of claim 35, wherein the plurality of computational elements represents a layer of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27; output of a neural network is an output layer of neurons)*.

Claim 37

Buck discloses the accelerator system of claim 36, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 38

Buck discloses the accelerator system of claim 21, wherein the memory is configured to store intermediate results of the sequence of computations *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)* and the controller is configured to transfer the output data from the memory to a main memory without transferring the intermediate results *(Buck: figure 5, temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W)*.

Claim 39

Buck discloses the accelerator system of claim 21, wherein the controller is configured to initiate transfer of at least a portion of the input data *(Buck: column 10, lines 36-43)* and to transfer the at least a portion of the output data in parallel *(figure 5, element 585; and column 6, lines 35-40, at least ISA bus is 8-bit wide indicating at least a portion of data is in parallel)*.

Claim 40

Buck discloses the accelerator system of claim 21, wherein the input data is acquired by at least one of a video camera, a microphone, or a cell recording electrode *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64)*.

Claim 42

Buck discloses a method of performing a sequence of computations representing an artificial neural network *(Buck: column 21, lines 61-64)*, the method comprising:

receiving, at a central processing unit (CPU), first input data acquired from an external system *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15; external systems ranging from input devices to the GPU)*;

performing, by a graphics processing unit (GPU), a first computation in the sequence of computations on first input data to generate first output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, computations in the sequence of computations representing respective layers of neurons in the artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*, an output of a computation in the sequence of computations representing an output of a neuron in the artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*;

storing, in a memory operably coupled to the GPU, the first input data and the first output data *(Buck: figure 5, element 230)*; and

transferring second input data into the memory during performance of the sequence of computations by the GPU *(Buck: figure 5, inputs and outputs; column 11, line 59 to column 12, line 6)*.

Claim 43

Buck discloses the method of claim 42, wherein transferring the second input data comprises transferring the second input data via a bus operably coupled to the central processing unit (CPU) *(Buck: figure 2)*.

Claim 47

Buck discloses the method of claim 42, further comprising: storing, in a first memory partition of the memory, parameters common to all of the computations in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Claim 48

Buck discloses the method of claim 47, further comprising: storing, in a second memory partition of the memory, data specific to the first computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Claim 49

Buck discloses the method of claim 48, further comprising: storing, in the second memory partition, external input data patterns, representations of internal variables, an input of the computation in the sequence of computations, and the output of the computation in the sequence of computations *(Buck: figure 5, element 530; column 10, lines 61-62; column 15, lines 49-60; column 10, lines 43-44; figure 5, element 510)*.

Claim 50

Buck discloses the method of claim 42, wherein storing the first output data comprises: accumulating, in the memory, outputs of computational elements executed by the GPU in performing the first computation in the sequence of computations *(Buck: column 11, lines 50-52)*.

Claim 51

Buck discloses the method of claim 42, further comprising: storing, in the memory, an output of a previous computation in the sequence of computations; and accessing, by the GPU, the output of the previous computation during performance of the computation in the sequence of computations *(Buck: figure 5, at least element 510; column 11, line 12 to column 12, line 15)*.

Claim 52

Buck discloses the method of claim 42, wherein performing the first computation comprises executing a plurality of computational elements representing a layer of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27; output of a neural network is an output layer of neurons)*.

Claim 53

Buck discloses the method of claim 52, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 54

Buck discloses the method of claim 42, further comprising: acquiring the input data with at least one of a video camera, a microphone, or a cell recording electrode *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64)*.

Claim 56

Buck discloses an accelerator system *(Buck: figure 2, element 190)* comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network *(Buck: column 21, lines 61-64)*, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*; and

a memory, operably coupled to the GPU, to store parameters common to all of the computations in the sequence of computations and data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*; and

a controller, operably coupled to the GPU and the memory, to transfer at least a portion of the input data into the memory in parallel with performance of the sequence of computations by the GPU, and to transfer at least a portion of the output data from the memory in parallel with performance of the sequence of computations by the GPU *(Buck: column 11, line 12 to column 12, line 15)*.

Claim 57

Buck discloses the accelerator system of claim 56, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 58

Buck discloses the accelerator system of claim 56, wherein the memory is configured to hold: external input data patterns *(Buck: figure 5, element 530; column 10, lines 61-62)*; representations of internal variables *(Buck: column 15, lines 49-60)*; an input to the computation in the sequence of computations *(Buck: column 10, lines 43-44)*; and the output of the computation in the sequence of computations *(Buck: figure 5, element 510)*.

Claim 59

Buck discloses the accelerator system of claim 56, wherein the memory is configured to accumulate the output of the computation in the sequence of computations during execution of the plurality of computational elements by the GPU *(Buck: column 11, lines 50-52)*.

Claim 60

Buck discloses the accelerator system of claim 56, wherein the memory is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to the plurality of computational elements during the computation in the sequence of computations *(Buck: figure 5, at least element 530; column 11, line 12 to column 12, line 15)*.

Claim 62

Buck discloses a method comprising:

receiving, at a central processing unit (CPU), input data acquired from an external system *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15; external systems ranging from input devices to the GPU)*;

performing, with a graphics processing unit (GPU), a sequence of computations on input data to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network *(Buck: column 21, lines 61-64)*, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*;

storing, in a memory operably coupled to the GPU, parameters common to all of the computations in the sequence of computations; in the memory, data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*; and

transferring the input data into the memory and at least a portion of the output data from the memory during performance of the sequence of computations by the GPU *(Buck: column 11, line 12 to column 12, line 15)*.

Claim 63-66

The limitations of claims 63-66 correspond to the limitations of claims 57-61 and as such are rejected in a corresponding manner.

Claim 68

Buck discloses a method of executing computations representing an artificial neural network *(Buck: column 21, lines 54-64)* on a computer system comprising at least one central processing unit (CPU) *(Buck: figure 2, element 120)*, a processing unit *(Buck: figure 2, element 220)*, a first memory partition *(Buck: figure 2, element 230; column 21, lines 46-49, multiple memory partitions/locations; see also column 15, lines 46-60 and figure 9)*, and a second memory partition *(Buck: figure 2, element 230; column 21, lines 46-49, multiple memory partitions/locations; see also column 15, lines 46-60 and figure 9)*, the method comprising:

executing, by the at least one CPU, a user interaction stream *(Buck: column 9, lines 35-36)*, the user interaction stream controlling transfer of inputs to the artificial neural network to the first memory partition and the second memory partition *(Buck: figures 2 and 5; column 10, lines 41-43; column 23, lines 13-20)*;

executing, by the processing unit, a computational stream *(Buck: column 2, lines 24-28; column 8, lines 39-44)*, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network *(Buck: figure 5, elements 220 and 545)*.

Claim 69

Buck discloses the method of claim 68, wherein the user interaction stream controls data exchange between the user interaction stream and the computational stream outside of execution of the computations representing the artificial neural network *(Buck: figure 2, element 130 separated by element 120)*.

Claim 70

Buck discloses the method of claim 68, wherein executing the user interaction stream comprises: controlling setting and editing of computational elements of the computations representing the artificial neural network *(Buck: figure 5, see at least 'P')*.

Claim 71

Buck discloses the method of claim 68, wherein executing the user interaction stream comprises: controlling setting and editing of parameters of the computations representing the artificial neural network *(Buck: figure 5, see at least 'W', 'L', 'ERRORS')*.

Claim 72

Buck discloses the method of claim 68, wherein executing the user interaction stream comprises: controlling setting and editing of parameters of the inputs to the artificial neural network *(Buck: figure 5, see at least 'X'' and 'T')*.

Claim 74

Buck discloses the method of claim 68, wherein executing the user interaction stream comprises: parsing elements to be used in the computations representing the artificial neural network *(Buck: column 2, lines 10-14; column 6, lines 10-16; column 10, lines 36-38; column 14, line 19 to column 15, line 18; programming and languages include parsing)*.

Claim 75

Buck discloses the method of claim 74, wherein the processing unit comprises a graphics processing unit (CPU) and executing the user interaction stream further comprises: converting the elements into GPU programs *(Buck: column 2, lines 10-14; column 6, lines 10-16; column 10, lines 36-38; column 14, line 19 to column 15, line 18; programming and languages include converting)*.

Claim 76

Buck discloses the method of claim 75, wherein executing the user interaction stream comprises: compiling the GPU programs *(Buck: column 2, lines 10-14; column 6, lines 10-16; column 10, lines 36-38; column 14, line 19 to column 15, line 18; DirectX includes compilation)*.

Claim 77

Buck discloses the method of claim 76, wherein executing the user interaction stream comprises: transferring the GPU programs to the second memory partition *(Buck: figure 5; and column 10, lines 36-41)*.

Claims 79-83 and 85-87

The limitations of claims 79-83 and 85-87 correspond to the limitations of claims 68-72 and 74-77 and as such are rejected in a corresponding manner.

Claim 88

Buck discloses the system of claim 87, wherein the user interaction stream transfers the GPU programs to at least one of the first memory partition or the second memory partition *(Buck: figure 5; column 10, lines 41-43; column 23, lines 13-20; figure 2, element 230; column 21, lines 46-49)*.

Claims 90, 92-94, 95, and 97-100

The limitations of claims 90, 92-94, 95, and 97-100 correspond to the limitations of claims 1-7, 10-16, and 19-20 and as such are rejected in a corresponding manner.

Claims 117-123, 126-132, and 135-136

The limitations of claims 117-123, 126-132, and 135-136 correspond to the limitations of claims 1-7, 10-16, 19-20, 42-43 and 47-54 and as such are rejected in a corresponding manner.

Claim 137

Buck discloses a system for executing an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*, the system comprising:

a first processing unit to receive input data in real time from a sensor *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15)*;

a memory, operably coupled to the first processing unit, to store a first portion of the input data *(Buck: figure 5, element 230)*;

a second processing unit, operably coupled to the memory, to generate first output data from the first portion of the input data, the first output data representing an output of at least one neuron in a first layer of neurons in the artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*;

a controller, operably coupled to the second processing unit and the memory, to store a second portion of the input data in the memory while the second processing unit generates the first output data *(Buck: figure 5, inputs and outputs; column 11, line 59 to column 12, line 6; column 23, lines 7-9)*.

Claim 138

Buck discloses the system of claim 137, wherein the sensor is a microphone *(Buck: column 7, lines 51-59)*.

Claim 140

Buck discloses the system of claim 137, wherein the second processing unit is further configured to generate second output data from the second portion of the input data, the second output data representing an output of at least one neuron in a second layer of neurons in the artificial neural network *(Buck: figure 5, inputs and outputs; column 11, line 59 to column 12, line 6; figure 7; column 21, line 65 to column 22, line 27)*.

Claim 141

Buck discloses the system of claim 140, further comprising: a random access memory, operably coupled to the first processing unit, to receive the first output data from the memory while the second processing unit generates the second output data *(Buck: figure 1, element 132; column 6, lines 53-59; column 7, lines 7-18; column 11, line 59 to column 12, line 6)*.

Claim 8-9, 17-18, 44-46, 96, 124-125, and 133-134 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,525,547 B1 *(herein **Diard**)*.

Claim 8

To the extent ***Buck*** does not explicitly state the computer system of claim 1, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one graphics processing unit has begun to perform another sequence of computations, ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 13, lines 4-25)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of ***Buck*** with output transferred from the accelerator memory to the main memory after the GPU has started other computations as suggested by ***Diard***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both ***Buck*** and ***Diard*** are directed to CPU and GPU processing.

Claim 9

To the extent ***Buck*** does not explicitly state the computer system of claim 8, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one graphics processing unit. ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 13, lines 4-25)*. ***Diard*** further demonstrates simultaneous transfer of input data and output data by an accelerator controller *(Diard: column 6, lines 10-26; column 13, lines 4-25)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of ***Buck*** with initiating transfer of input data and transferring output data in parallel with a GPU's computation in another sequence of computations as suggested by ***Diard***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both ***Buck*** and ***Diard*** are directed to CPU and GPU processing.

Claims 17-18

The limitations of claims 17-18 correspond to the limitations of claims 8-9 and as such are rejected in a corresponding manner.

Claim 44

Buck discloses the method of claim 42, further comprising: transferring the first output data from the memory to another memory during a second computation in the sequence of computations. ***Diard*** demonstrates that it was known at the time of invention to transfer output from a memory to another

memory during a computation *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 13, lines 4-25)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of ***Buck*** with output transferred from the accelerator memory to the main memory after the GPU has started other computations as suggested by ***Diard***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both ***Buck*** and ***Diard*** are directed to CPU and GPU processing.

Claim 45

Buck discloses the method of claim 44, further comprising: storing intermediate results of the sequence of computations in the memory, and wherein transferring the first output data from the memory to the other memory occurs without transferring the intermediate results of the sequence of computations *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W; figure 5, temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W)*.

Claim 46

Buck discloses the method of claim 44, wherein transferring the second input data and transferring the first output data occurs in parallel *(Buck: figure 5, element 585; and column 6, lines 35-40, at least ISA bus is 8-bit wide indicating at least a portion of data is in parallel)*. Also, ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3;*

column 13, lines 4-25). ***Diard*** further demonstrates simultaneous transfer of input data and output data by an accelerator controller *(Diard: column 6, lines 10-26; column 13, lines 4-25)*.

Claim 96

To the extent ***Buck*** does not explicitly state the system of claim 95, wherein the first memory partition and the second memory partition are logical partitions, ***Diard*** demonstrates that it was known at the time of invention for various memory to be shared or logically partitioned between graphics units and CPUs *(Diard: column 4, lines 5-11)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the memory of ***Buck*** as logical partitions as suggested by ***Diard***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both ***Buck*** and ***Diard*** are directed to CPU and GPU processing.

Claims 124-125 and 133-134

The limitations of claims 124-125 and 133-134 correspond to the limitations of claims 8-9 and 17-18 and as such are rejected in a corresponding manner.

Claims 41, 55, and 91 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,119,810 B2 *(herein **Sumanaweera**)*.

Claim 41

To the extent ***Buck*** does not explicitly state the accelerator system of claim 21, wherein the input data is loaded from disk. ***Sumanaweera*** demonstrates that it was known at the time of invention to load GPUs

from disk *(Sumanaweera: column 1, lines 10-19)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement ***Buck*** with input data loaded from disk as suggested by ***Sumanaweera***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claims 55 and 91

The limitations of claims 55 and 91 correspond to the limitations of claim 41 and as such are rejected in a corresponding manner.

Claims 73, 78, 84, 89, and 139 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 2006/0129506 A1 *(herein **Edelman**)*.

Claim 73

To the extent, ***Buck*** does not explicitly state the method of claim 68, wherein executing the user interaction stream comprises: specifying an output to be saved to disk and/or displayed on a screen. ***Edelman*** demonstrates that it was known at the time of invention to save computations of an output stream to a disk *(Edelman: paragraphs 0056, and 0079)* and for the output stream to be related to an artificial or simulated neural system *(Edelman: paragraphs 0056, and 0079)*. Further, ***Buck*** demonstrates CPU controlled outputs and disk memory disk *(Buck: figures 1 and 2; column 3, lines 46-51; column 9, lines 15-37)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the artificial neural network of ***Buck*** with saving output to disk as suggested by ***Edelman***'s teachings. This implementation would have been obvious because one of ordinary skill in the

art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 78

To the extent ***Buck*** does not state the method of claim 68, further comprising: executing, by the at least one CPU, a data output stream, the data output stream controlling transfer of outputs of the computations representing the artificial neural network to disk. ***Edelman*** demonstrates that it was known at the time of invention to save computations of an output stream to a disk *(Edelman: paragraphs 0056, and 0079)* and for the output stream to be related to an artificial or simulated neural system *(Edelman: paragraphs 0056, and 0079)*. Further, ***Buck*** demonstrates CPU controlled outputs and disk memory disk *(Buck: figures 1 and 2; column 3, lines 46-51; column 9, lines 15-37)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the artificial neural network of ***Buck*** with saving output to disk as suggested by ***Edelman***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claims 84 and 89

The limitations of claims 84 and 89 correspond to the limitations of claims 73 and 78 and as such are rejected in a corresponding manner.

Claim 139

Buck does not explicitly state wherein the sensor is a video camera. ***Edelman*** demonstrates that it was known at the time of invention to make use of video camera input sensors *(Edelman: paragraph 0049)*

in an environment related to artificial neural networks *(Edelman: paragraph 0006)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement an input sensor of *Buck* with or as a video camera as suggested by *Edelman*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both references are directed to artificial neural networks.

Claims 101-104 and 106-116 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2001/0010034 A1 *(herein **Burton**)* in view of US 7,219,085 B2 *(herein **Buck**)*.

Claim 101

Burton discloses a method for performing a numerical simulation over a plurality of computational cycles including at least a first computational cycle and a second computational cycle *(Burton: figure 7; paragraph 0075)*, the method comprising:

storing, in a first partition of a memory, first output data during the second computational cycle, the first output data generated during the first computational cycle *(Burton: figure 7; paragraph 0075; the current cycle, which corresponds to the claimed second cycle, uses data that it is storing in memory area A and which was previously generated in a first cycle)* and being referenced by a first pointer *(Burton: figure 7, element 38)*;

performing, by a processing unit operably coupled to the memory, at least one calculation on the first output data during the second computational cycle, the at least one calculation yielding second output data *(Burton: figure 7, elements 32-36; paragraph 0075; the current cycle generates new output data)*;

storing, in a second partition of the memory, the second output data, the second output data being referenced by a second pointer *(Burton: figure 7, elements 36 and 38; paragraph 0075; writing output to memory area B)*; and

swapping the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles *(Burton: figure 7, element 38; the new output data can now be used for another future cycle, which corresponds to the claimed third cycle)*.

Burton does not explicitly state: "a numerical simulation representing an artificial neural network"; "a first computational cycle representing a first layer of neurons in the artificial neural network"; "a second computational cycle representing a second layer of neurons in the artificial neural network"; "the first output data representing at least one output of the first layer of neurons"; and "the second output data representing at least one output of the second layer of neurons". As noted above, *Burton* discloses a numerical simulation with first and second computational cycles. ***Buck*** demonstrates that it was known at the time of invention to perform numerical simulations representing artificial neural networks *(Buck: column 10, lines 45-46; column 11, lines 42-43; column 21, lines 60-64)*; to represent layers of the artificial neural network as a group of computations or a computational cycle *(Buck: figure 7; column 21, line 60 to column 22, line 27)*; and to pass outputs from one layer as inputs to another layer *(Buck: column 22, lines 1-7)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the first and second cycles of *Burton*'s numerical simulation as representing layers of an artificial neural network, and thus the first and second cycles generating a respective first and second output of a respective layer of neurons, as suggested by *Buck*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another

yielding a predictable result using an acceptable piece of prior art; and both are directed to simulated environments.

Claim 102

Burton discloses the method of claim 101, further comprising: transferring input data into the memory from another memory coupled to a central processing unit (CPU) before the first computational cycle *(Burton: figures 6-7; paragraphs 0074-0075; input data for the claim labeled "first cycle" coming from an even earlier cycle of Burton's loop)*.

Claim 103

Burton discloses the method of claim 101, further comprising: providing, to a central processing unit (CPU), first output data generated during the first computational cycle after a conclusion of the first computational cycle *(Burton: figures 6-7; paragraphs 0074-0075; showing the CPU is the controlling/executing everything and is thus provided with the output and adjusting/swapping the pointers in all cycles)*.

Claim 104

Burton discloses the method of claim 101, further comprising: transferring the first output data from the memory to another memory during the second computational cycle *(Burton: figure 7, element 38)*.

Claim 106

Burton discloses the method of claim 101, further comprising: dictating, to the processing unit, an order of execution of instructions *(Burton: paragraph 0075, "compile time scheduling")*.

Claim 107

Burton discloses the method of claim 101, wherein performing the at least one calculation comprises receiving, by the processing unit, at least the portion of the input data and a series of instructions and executing, by the processing unit, the instructions on the at least the portion of the input data *(Burton: paragraph 0074-0075, 0078, and 0042; executing compiled instructions from memory describing the simulation and the hardware, and including data of the simulation)*.

Claim 108

Burton discloses a method of performing a numerical simulation *(Burton: figure 7; paragraph 0075)*, the method comprising:

performing, by a processing unit, at least one computation during a first computational cycle of the numerical simulation *(Burton: figure 7; paragraph 0075; first iteration of the loop)*;

storing, by a memory operably coupled to the processing unit, data used to perform the at least one computation *(Burton: figure 7; paragraph 0075; memory area A storing inputs to the first iteration of the loop)*;

controlling performance of the at least one computation by the processing unit *(Burton: figure 7; paragraph 0075; executing generating output during the first iteration of the loop)*;

controlling transfer of the output data from the processing unit to the memory during the first computational cycle *(Burton: figure 7; paragraph 0075; writing output to memory area B during the first iteration of the loop)*; and

controlling transfer of input data to the memory for use by the processing unit during a second computational cycle of the numerical simulation *(Burton: figure 7; paragraph 0075; swapping pointers and then reading from memory area A the old output, which is now new input, during the second iteration of the loop)*.

Burton does not explicitly state: "a numerical simulation representing an artificial neural network"; "a first computational cycle representing a first layer of neurons in the artificial neural network"; and "a second computational cycle representing a second layer of neurons in the artificial neural network". As noted above, *Burton* discloses a numerical simulation with first and second computational cycles. ***Buck*** demonstrates that it was known at the time of invention to perform numerical simulations representing artificial neural networks *(Buck: column 10, lines 45-46; column 11, lines 42-43; column 21, lines 60-64)*; to represent layers of the artificial neural network as a group of computations or a computational cycle *(Buck: figure 7; column 21, line 60 to column 22, line 27)*; and to pass outputs from one layer as inputs to another layer *(Buck: column 22, lines 1-7)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the first and second cycles of *Burton*'s numerical simulation as representing layers of an artificial neural network, and thus the first and second cycles generating a respective first and second output of a respective layer of neurons, as suggested by *Buck*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to simulated environments.

Claim 109

Burton discloses the method of claim 108, further comprising: dictating an order of execution of instructions to the processing unit *(Burton: paragraph 0075, "compile time scheduling")*.

Claim 111

Burton discloses the method of claim 108, further comprising: perform successive computational cycles of the numerical simulation by feeding output data generated by the processing unit from a previous computational cycle and input data for a next computational cycle into the processing unit *(Burton: figure 7; paragraph 0075; output of the preceding cycle is input for the next cycle)*.

Claim 112

Burton discloses a system for performing a numerical simulation on input data *(Burton: figure 7; paragraph 0075)*, the system comprising:

a central processing unit (CPU) to provide first input data *(Burton: figure 6, element 22)*;

a memory, operably coupled to the CPU, to store the first input data in a first partition, referenced by first pointer, before a first computational cycle of the numerical simulation *(Burton: figure 6, element 24)*;

a processing unit, operably coupled to the memory, to perform, during a first computational cycle, at least one calculation on the first portion of the input data so as to generate first output data *(Burton: figure 6, CPU; figure 7; paragraph 0075; first iteration of the loop)*; and

a controller, operably coupled to the processing unit and the memory, to store the first output data in a second partition of the memory *(Burton: figure 7; paragraph 0075; memory area B)*, the second partition referenced by a second pointer, and to swap the first pointer with the second pointer at the end of the first computational cycle, such that the first output data becomes an input for a second computational cycle of the numerical simulation *(Burton: figure 7, element 38; paragraph 0075)*.

Burton does not explicitly state: "artificial neural network"; "a first layer of neurons in the artificial neural network"; and "a second layer of neurons in the artificial neural network". As noted above, *Burton* discloses a numerical simulation with first and second computational cycles. ***Buck*** demonstrates

that it was known at the time of invention to perform numerical simulations representing artificial neural networks *(Buck: column 10, lines 45-46; column 11, lines 42-43; column 21, lines 60-64)*; to represent layers of the artificial neural network as a group of computations or a computational cycle *(Buck: figure 7; column 21, line 60 to column 22, line 27)*; and to pass outputs from one layer as inputs to another layer *(Buck: column 22, lines 1-7)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the first and second cycles of *Burton*'s numerical simulation as representing layers of an artificial neural network, and thus the first and second cycles generating a respective first and second output of a respective layer of neurons, as suggested by *Buck*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to simulated environments.

Claim 110

Burton does not explicitly state the method of claim 108, further comprising: transferring output data from the memory to a main memory while the processing unit performs the at least one computation. ***Buck*** demonstrates that it was known at the time of invention to transfer output data to another memory, for example out to a main memory of a CPU, while the processing unit performs a computation *(Buck: column 11, line 12 to column 12, line 26; also figures 1 and 2)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the processing unit of ***Burton*** as a CPU making use of a GPU with transferring output data to main memory during a computation as suggested by ***Buck****'s* teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application

of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 113

Burton does not explicitly state the system of claim 112, wherein the processing unit comprises a graphics processing unit. ***Buck*** demonstrates that it was known at the time of invention to off load processing to a graphics processing unit *(Buck: figure 5)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the processing unit of ***Burton*** as a CPU making use of a GPU as suggested by ***Buck's*** teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 114

Burton does not explicitly state the system of claim 112, wherein the controller is configured to send instructions for performing the at least one calculation to the processing unit. ***Buck*** demonstrates that it was known at the time of invention for a controller to send instructions for performing a calculation to a processing unit *(Buck: column 10, lines 34 to column 12, line 26)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the processing unit of ***Burton*** as a GPU with a controller to send instructions for performing a calculation as suggested by ***Buck's*** teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 115

Burton does not explicitly state the system of claim 112, wherein the controller is further configured to transfer the first output data to another memory during computation of the second layer of neurons. ***Buck*** demonstrates that it was known at the time of invention to transfer output data to another memory, for example out to a CPU, during a second computational cycle *(Buck: column 11, line 12 to column 12, line 26)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the processing unit of ***Burton*** as a CPU making use of a GPU with a controller transferring output data to another memory during a second computation cycle as suggested by ***Buck's*** teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 116

Burton does not explicitly state the system of claim 112, wherein the memory further comprises: a third partition to store internal variables; and a fourth partition to store data used as input at a particular layer of neurons of the artificial neural network. ***Buck*** demonstrates that it was known at the time of invention for a memory to have a partition to store internal variables and a partition to store data used as input at a computational cycle *(Buck: figure 5, element 230)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the processing unit of ***Burton*** as a CPU making use of a GPU with a memory partitioned with partitions for internal variables, data used as inputs to a computational cycle, as well as input data and output data as suggested by ***Buck's*** teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 105 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2001/0010034 A1 *(herein **Burton**)* in view of US 7,219,085 B2 *(herein **Buck**)* in further view of US Patent 5,142,665 *(herein **Bigus**)*.

Claim 105

Burton does not explicitly state the method of claim 101, further comprising: displaying, with a video system, an indication of the numerical simulation. ***Bigus*** demonstrates that it was known at the time of invention to display simulation results *(Bigus: column 3, lines 32-34; column 6, lines 35-41; column 10, lines 29-33)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the simulations of ***Burton*** with displaying an indication of the simulation as suggested by ***Bigus'*** teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Response to Arguments

Patent Owner's arguments filed 03/20/2019 (herein: Remarks) have been fully considered but they are not persuasive. Patent Owner argues: 1) various claims do not invoke 251 recapture *(Remarks: pages 28-40)*; 2) various claims are not anticipated under 35 USC 102 by *Buck (Remarks: pages 40-49)*; and 3) various claims are not anticipated by *Burton (Remarks: pages 49-50)*.

First, Patent Owner's arguments with regard to 251 recapture are addressed as follows and taken in consideration with the above recapture analysis: 1) for claim 21, the arguments are persuasive for the claim is at least materially narrowed; 2) for claim 42, the arguments are persuasive for the claim

is at least materially narrowed; 3) for claim 56, the arguments are persuasive for the claim is at least materially narrowed; 4) for claim 62, the arguments are persuasive for the claim is at least materially narrowed; 5) for claim 68, the arguments are persuasive for the claim is at least materially narrowed; 6) for claim 79, the arguments are persuasive for the claim is at least materially narrowed; 7) for claim 90, the arguments are persuasive for the claim is at least materially narrowed; 8) for claim 95, the arguments are persuasive for the claim is at least materially narrowed; 9) for claim 101, the arguments are persuasive for the claim is at least materially narrowed; 10) for claim 108, the arguments are persuasive for the claim is at least materially narrowed; and 11) for claim 112, the arguments are persuasive for the claim is at least materially narrowed.

Second, with regard to *Buck* and in particular to claim 1, Patent Owner argues "Buck is silent on transferring data to a GPU or GPU memory while the GPU is performing a sequence of computations" *(Remarks: page 45, section 9.2)*. The claim language requires, a controller to transfer at least a portion of data "during performance of the sequence of computations". The term "sequence of computations" is not defined in the originally filed disclosure. Under the broadest reasonable interpretation of the claim language by one of ordinary skill in the art at the time of invention, "sequence of computations" includes every computation (or even operation) performed by the training phase in the GPU of *Buck* as cited in the rejection *(Buck: column 11, line 12 to column 12, line 36)*. Note, "the CPU can get feedback from the GPU inside the training loop … the GPU send the CPU the updated trainable parameters (W) and the errors" *(Buck: column 11, lines 60-64)*. At least a portion of data is therefore transmitted "during performance of the sequence of computations". Additionally, *Buck* shows parameters and errors being transmitted to the CPU from the GPU *(Buck: figure 5, element 565)* and other data being transmitted to the GPU from the CPU *(Buck: figure 5, element 570)*. This happens before the training is completed and

therefore at least a portion of input and output data is transferred "during performance of the sequence of computations". Patent Owner's arguments are not persuasive.

With regard to *Buck* and in particular claim 3, Patent Owner argues the cited prior art does not "say anything about the rate at which the CPU receives input data". However, *Buck* does indicate a GPU operates at a rate/speed that is faster than a CPU *(Buck: column 14, lines 40-44, "algorithms run an order of magnitude faster on a GPU than on a CPU")*. Therefore, the processing of the CPU, including the receiving of input data, is at a rate different than the processing of the GPU. Further, *Buck* gives an example implementation using a CPU processor and a GPU processor with differing speeds *(Buck: column 21, lines 43-53)*. Patent Owner's arguments are not persuasive.

With regard to *Buck* and in particular claim 11, Patent Owner argues the cited prior art does not "say anything about acquiring input data in real time with a video camera, microphone or cell recording electrode". The claim language requires, "at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time". The term "real time" is not defined in the originally filed disclosure. Under the broadest reasonable interpretation of the claim language by one of ordinary skill in the art at the time of invention, "real time" operations are those in which the machine's activities match the human perception of time or those in which computer operations proceed at the same rate as a physical or external process. *Buck* discloses microphone *(Buck: column 7, lines 56-57)*. *Buck* also discloses "in interactive use applications, such as speech recognition" *(Buck: column 9, lines 33-36)*. A microphone matches the human perception of a machine acquiring input as it happens or in real time. Further, an interactive application matches the human perception of a machine acquiring input as it happens or in real time. Also, note the claim requires the input device, such as a microphone. The claim language, "to acquire the input data in real time" is intended use and does not add significantly to the claim language as microphones record as events are happening, or in other words in real time.

With regard to *Buck* and in particular claim 12, Patent Owner argues *Buck* "in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data" *(Remarks: pages 46-47, section 9.5)*. Patent Owner further suggests the transfer is "interleaved", presumably to indicate an alternating pattern of transmitting and computing. However, as discussed previously for claim 1, *Buck* shows transferring at least a portion of input and output data once the "sequence of computations" have begun and before they have finished. The claim language requires "in parallel with performing" the sequence of computations, not in parallel with any one particular computation. Therefore, the arguments are not persuasive.

With regard to *Buck* and in particular claims 21, 42, 56, 62, 90, 95, 117, and 128, Patent Owner argues the cited prior art does not disclose variations of transferring data during performance of a sequence of computations. These arguments are not persuasive for the same reasons as discussed above for claim 1.

With regard to *Buck* and in particular claims 68 and 79, Patent Owner argues *Buck* does not disclose a "user interaction stream" because "the data transferred from the GPU to the CPU [of *Buck*] is not an input to an artificial neural network". Due to the amendment, the rejections of claims 68 and 79 has been altered and this argument is no longer applicable.

Third, with regard to claims 101, 108, and 112 and rejection under *Burton*, the arguments are not persuasive. Upon review of the amended claims, the rejections are now rejected over *Burton* in view of *Buck*. *Buck* does show a neural networks. The arguments are unpersuasive.

Finally, the obviousness rejections and the dependent claim rejections are maintained as indicated and based upon the above analysis of Patent Owner's arguments of 03/20/2019. The previous 112 rejections are withdrawn in view of the Patent Owner's response of 03/20/2019. The objections and

rejections concerning the Declaration are withdrawn in view of the Patent Owner's response of 03/20/2019. Having addressed all the substantive issues, any objections and rejections are maintained as above indicated.

Conclusion

Patent Owner's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Patent Owner is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Correspondence Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to William H. Wood whose telephone number is (571)272-3736 and fax number is (571)273-3736. The examiner can normally be reached on Monday-Friday 7:30am-4pm. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Alexander Kosowski can be reached on (571)272-3744. The Central Reexamination Unit's (CRU) fax number is (571)273-9900.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR systems, see https://ppair.uspto.gov/epatent/portal/home. For questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/William H. Wood/
Primary Examiner, CRU 3992

Conferees:

/RSD/

/ALEXANDER J KOSOWSKI/
Supervisory Patent Examiner, Art Unit 3992

Applicant-Initiated Interview Summary	**Application No.** 15/808,201	**Applicant(s)** Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	**Art Unit** 3992	**AIA (FITF) Status** No

All participants (applicant, applicants representative, PTO personnel):

(1) WILLIAM H. WOOD.

(2) Max Colice.

(3) Alexander Kosowski.

(4) Adam Basehoar.

Date of Interview: 08 February 2019.

Type: ☐ Telephonic ☑ Video Conference
☐ Personal [copy given to: ☐ applicant ☐ applicant's representative]

Exhibit shown or demonstration conducted: ☑ Yes ☐ No.
If Yes, brief description: Interview Agenda (attached).

Issues Discussed ☐101 ☐112 ☑102 ☐103 ☑Others
(For each of the checked box(es) above, please describe below the issue and detailed description of the discussion)

Claim(s) discussed: 1,12,21,42,56,62,68,79,90,95,101,108,112,117 and 128.

Identification of prior art discussed: Buck.

Substance of Interview
(For each issue discussed, provide a detailed description and indicate if agreement was reached. Some topics may include: identification or clarification of a reference or a portion thereof, claim interpretation, proposed amendments, arguments of any applied references etc...)

Also attending: Andrew Baluch. Briefly discussed: submitting missing IDs references (all agreed); and more fully explaining the error statement in the Declaration, with specific claim language and a direct mention of at least one broadened claim (all agreed). Further discussed: application of prior art Buck to the claim language, including the concept of a "controller" and the timing of the sequence of computations relative to transfers of input and output data; and various recapture rejections and possible solutions, including possible narrowing limitations and possible overlooked aspects or embodiments. All parties agreed understanding of the invention and rejections was clarified and advanced..

Applicant recordation instructions: The formal written reply to the last Office action must include the substance of the interview. (See MPEP section 713.04). If a reply to the last Office action has already been filed, applicant is given a non-extendable period of the longer of one month or thirty days from this interview date, or the mailing date of this interview summary form, whichever is later, to file a statement of the substance of the interview.

Examiner recordation instructions: Examiners must summarize the substance of any interview of record. A complete and proper recordation of the substance of an interview should include the items listed in MPEP 713.04 for complete and proper recordation including the identification of the general thrust of each argument or issue discussed, a general indication of any other pertinent matters discussed regarding patentability and the general results or outcome of the interview, to include an indication as to whether or not agreement was reached on the issues raised.

☑ Attachment

/William H. Wood/ Primary Examiner, Art Unit 3992	

Summary of Record of Interview Requirements

Manual of Patent Examining Procedure (MPEP), Section 713.04, Substance of Interview Must be Made of Record
A complete written statement as to the substance of any face-to-face, video conference, or telephone interview with regard to an application must be made of record in the application whether or not an agreement with the examiner was reached at the interview.

Title 37 Code of Federal Regulations (CFR) 1.133 Interviews
Paragraph (b)

In every instance where reconsideration is requested in view of an interview with an examiner, a complete written statement of the reasons presented at the interview as warranting favorable action must be filed by the applicant. An interview does not remove the necessity for reply to Office action as specified in §§ 1.111, 1.135. (35 U.S.C. 132)

37 CFR §1.2 Business to be transacted in writing.
All business with the Patent or Trademark Office should be transacted in writing. The personal attendance of applicants or their attorneys or agents at the Patent and Trademark Office is unnecessary. The action of the Patent and Trademark Office will be based exclusively on the written record in the Office. No attention will be paid to any alleged oral promise, stipulation, or understanding in relation to which there is disagreement or doubt.

The action of the Patent and Trademark Office cannot be based exclusively on the written record in the Office if that record is itself incomplete through the failure to record the substance of interviews.

It is the responsibility of the applicant or the attorney or agent to make the substance of an interview of record in the application file, unless the examiner indicates he or she will do so. It is the examiners responsibility to see that such a record is made and to correct material inaccuracies which bear directly on the question of patentability.

Examiners must complete an Interview Summary Form for each interview held where a matter of substance has been discussed during the interview by checking the appropriate boxes and filling in the blanks. Discussions regarding only procedural matters, directed solely to restriction requirements for which interview recordation is otherwise provided for in Section 812.01 of the Manual of Patent Examining Procedure, or pointing out typographical errors or unreadable script in Office actions or the like, are excluded from the interview recordation procedures below. Where the substance of an interview is completely recorded in an Examiners Amendment, no separate Interview Summary Record is required.

The Interview Summary Form shall be given an appropriate Paper No., placed in the right hand portion of the file, and listed on the "Contents" section of the file wrapper. In a personal interview, a duplicate of the Form is given to the applicant (or attorney or agent) at the conclusion of the interview. In the case of a telephone or video-conference interview, the copy is mailed to the applicants correspondence address either with or prior to the next official communication. If additional correspondence from the examiner is not likely before an allowance or if other circumstances dictate, the Form should be mailed promptly after the interview rather than with the next official communication.

The Form provides for recordation of the following information:
- Application Number (Series Code and Serial Number)
- Name of applicant
- Name of examiner
- Date of interview
- Type of interview (telephonic, video-conference, or personal)
- Name of participant(s) (applicant, attorney or agent, examiner, other PTO personnel, etc.)
- An indication whether or not an exhibit was shown or a demonstration conducted
- An identification of the specific prior art discussed
- An indication whether an agreement was reached and if so, a description of the general nature of the agreement (may be by attachment of a copy of amendments or claims agreed as being allowable). Note: Agreement as to allowability is tentative and does not restrict further action by the examiner to the contrary.
- The signature of the examiner who conducted the interview (if Form is not an attachment to a signed Office action)

It is desirable that the examiner orally remind the applicant of his or her obligation to record the substance of the interview of each case. It should be noted, however, that the Interview Summary Form will not normally be considered a complete and proper recordation of the interview unless it includes, or is supplemented by the applicant or the examiner to include, all of the applicable items required below concerning the substance of the interview.

A complete and proper recordation of the substance of any interview should include at least the following applicable items:
1) A brief description of the nature of any exhibit shown or any demonstration conducted,-
2) an identification of the claims discussed,
3) an identification of the specific prior art discussed,
4) an identification of the principal proposed amendments of a substantive nature discussed, unless these are already described on the Interview Summary Form completed by the Examiner,
5) a brief identification of the general thrust of the principal arguments presented to the examiner,
(The identification of arguments need not be lengthy or elaborate. A verbatim or highly detailed description of the arguments is not required. The identification of the arguments is sufficient if the general nature or thrust of the principal arguments made to the examiner can be understood in the context of the application file. Of course, the applicant may desire to emphasize and fully describe those arguments which he or she feels were or might be persuasive to the examiner.)
6) a general indication of any other pertinent matters discussed, and
7) if appropriate, the general results or outcome of the interview unless already described in the Interview Summary Form completed by the examiner.

Examiners are expected to carefully review the applicants record of the substance of an interview. If the record is not complete and accurate, the examiner will give the applicant an extendable one month time period to correct the record.

Examiner to Check for Accuracy

If the claims are allowable for other reasons of record, the examiner should send a letter setting forth the examiners version of the statement attributed to him or her. If the record is complete and accurate, the examiner should place the indication, Interview Record OK on the paper recording the substance of the interview along with the date and the examiners initials.

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019							
	1	✓	✓							
	2	✓	✓							
	3	✓	✓							
	4	✓	✓							
	5	✓	✓							
	6	✓	✓							
	7	✓	✓							
	8	✓	✓							
	9	✓	✓							
	10	✓	✓							
	11	✓	✓							
	12	✓	✓							
	13	✓	✓							
	14	✓	✓							
	15	✓	✓							
	16	✓	✓							
	17	✓	✓							
	18	✓	✓							
	19	✓	✓							
	20	✓	✓							
	21	✓	✓							
	22	✓	✓							
	23	✓	✓							
	24	✓	✓							
	25	✓	✓							
	26	✓	✓							
	27	✓	✓							
	28	✓	✓							
	29	✓	✓							
	30	✓	✓							
	31	✓	✓							
	32	✓	✓							
	33	✓	✓							
	34	✓	✓							
	35	✓	✓							
	36	✓	✓							
	37	✓	✓							
	38	✓	✓							
	39	✓	✓							
	40	✓	✓							
	41	✓	✓							
	42	✓	✓							

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019							
	43	✓	✓							
	44	✓	✓							
	45	✓	✓							
	46	✓	✓							
	47	✓	✓							
	48	✓	✓							
	49	✓	✓							
	50	✓	✓							
	51	✓	✓							
	52	✓	✓							
	53	✓	✓							
	54	✓	✓							
	55	✓	✓							
	56	✓	✓							
	57	✓	✓							
	58	✓	✓							
	59	✓	✓							
	60	✓	✓							
	61	✓	-							
	62	✓	✓							
	63	✓	✓							
	64	✓	✓							
	65	✓	✓							
	66	✓	✓							
	67	✓	-							
	68	✓	✓							
	69	✓	✓							
	70	✓	✓							
	71	✓	✓							
	72	✓	✓							
	73	✓	✓							
	74	✓	✓							
	75	✓	✓							
	76	✓	✓							
	77	✓	✓							
	78	✓	✓							
	79	✓	✓							
	80	✓	✓							
	81	✓	✓							
	82	✓	✓							
	83	✓	✓							
	84	✓	✓							
	85	✓	✓							
	86	✓	✓							
	87	✓	✓							
	88	✓	✓							
	89	✓	✓							
	90	✓	✓							
	91	✓	✓							
	92	✓	✓							
	93	✓	✓							
	94	✓	✓							

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019							
	95	✓	✓							
	96	✓	✓							
	97	✓	✓							
	98	✓	✓							
	99	✓	✓							
	100	✓	✓							
	101	✓	✓							
	102	✓	✓							
	103	✓	✓							
	104	✓	✓							
	105	✓	✓							
	106	✓	✓							
	107	✓	✓							
	108	✓	✓							
	109	✓	✓							
	110	✓	✓							
	111	✓	✓							
	112	✓	✓							
	113	✓	✓							
	114	✓	✓							
	115	✓	✓							
	116	✓	✓							
	117	✓	✓							
	118	✓	✓							
	119	✓	✓							
	120	✓	✓							
	121	✓	✓							
	122	✓	✓							
	123	✓	✓							
	124	✓	✓							
	125	✓	✓							
	126	✓	✓							
	127	✓	✓							
	128	✓	✓							
	129	✓	✓							
	130	✓	✓							
	131	✓	✓							
	132	✓	✓							
	133	✓	✓							
	134	✓	✓							
	135	✓	✓							
	136	✓	✓							
	137		✓							
	138		✓							
	139		✓							
	140		✓							
	141		✓							

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

CPC - Searched*		
Symbol	Date	Examiner
G06F 9/5027	10/10/2018	/WHW/
G06T 1/20	10/10/2018	/WHW/
G06T 1/60	10/10/2018	/WHW/
G06N 99/005	10/10/2018	/WHW/
G06N 3/063	10/10/2018	/WHW/
G06F 2209/509	10/10/2018	/WHW/

CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

Search Notes		
Search Notes	Date	Examiner
Prosecution History Review	10/10/2018	/WHW/
IDS Review	10/10/2018	/WHW/
Litigation Search	10/10/2018	/WHW/
EAST BRS Keyword Search	10/10/2018	/WHW/
NPL IEEE Search	10/10/2018	/WHW/
IDS Review	05/03/2019	/WHW/
EAST BRS	05/03/2019	/WHW/

/William H. Wood/ Primary Examiner, Art Unit 3992	

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

Interference Search			
US Class/CPC Symbol	US Subclass/CPC Group	Date	Examiner

/William H. Wood/ Primary Examiner, Art Unit 3992	

Doc code: IDS
Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (02-18)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	10083523		2018-09-25	Versace et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	15808201
	Filing Date	2017-11-09
	First Named Inventor	Anatoli Gorchechnikov
	Art Unit	3992
	Examiner Name	William H WOOD
	Attorney Docket Number	NRLA-001US03

	~~1~~	~~Brain inspired computing. Invited keynote address, Bionetics 2010, Boston, MA, USA. 1 page.~~
	2	Livitz G., Versace M., Gorchetchnikov A., Vasilkoski Z., Ames H., Chandler B., Leveille J. and Mingolla E. (2011) Adaptive, brain-like systems give robots complex behaviors, The Neuromorphic Engineer, : 10.2417/1201101.003500 February 2011. 3 pages.
	3	Salakhutdinov, R., & Hinton, G. E. (2009). Deep boltzmann machines. In International Conference on Artificial Intelligence and Statistics (pp. 448-455).
	4	SHERBAKOV, L. et al. 2012. CogEye: from active vision to context identification, youtube, retrieved from the Internet on 10-10-2017: URL://www.youtube.com/watch?v=i5PQk962B1k, 1 page.
	5	SHERBAKOV, L. et al. 2013. CogEye: system diagram module brain area function algorithm approx # neurons, retrieved from the Internet on 10-12-2017: URL://http://www-labsticc.univ-ubs.fr/~coussy/neucomp2013/index_fichiers/material/posters/NeuComp2013_final56x36.pdf, 1 page.
	6	Snider, Greg, et al. "From synapses to circuitry: Using memristive memory to explore the electronic brain." IEEE Computer, Volume 44(2). (2011): 21-28.
	7	VERSACE, TEDx Fulbright, Invited talk, Washington DC, 4/5/2014. 30 pages.

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/William H. Wood/	Date Considered	04/18/2019

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Christopher Max Colice/	Date (YYYY-MM-DD)	2019-03-20
Name/Print	Christopher Max Colice	Registration Number	65634

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S129	1	("10083523").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/04/18 14:51
S130	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/04/18 15:16
S131	1	S130 and sequence	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/04/18 15:16
S132	1	S130 and sequential	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/04/18 15:16
S133	1	S130 and real	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/04/18 16:10
S135	1	("7219085").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/04/24 15:13
S136	0	S135 and camera	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/04/24 15:14
S137	1	S135 and video	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/04/24 15:14
S138	45	(US-20170102950-$ or US-20070124363-$ or US-20030152288-$ or US-20080238917-$ or US-20080001957-$ or US-20040207618-$ or US-20020190994-$ or US-20070052713-$ or US-20060106591-$ or US-20110035737-$ or	US-PGPUB; USPAT; FPRS; DERWENT	OR	ON	2019/04/24 15:24

NVIDIA Ex1002

		US-20090222572-$ or US-20090027402-$ or US-20090027383-$ or US-20080181081-$ or US-20060233525-$ or US-20060114254-$ or US-20010052123-$ or US-20070165035-$ or US-20010010034-$ or US-20060129506-$).did. or (US-7119810-$ or US-7746350-$ or US-9349201-$ or US-8194085-$ or US-7737984-$ or US-7068278-$ or US-7522167-$ or US-6999603-$ or US-7526634-$ or US-7525547-$ or US-6367122-$ or US-6146277-$ or US-7219085-$ or US-7477256-$ or US-6952827-$ or US-5977977-$ or US-5099447-$ or US-5987443-$ or US-4994982-$ or US-5142665-$).did. or (WO-9316438-$ or WO-0227658-$).did. or (US-20050125369-$ or US-20050128206-$).did.				
S139	6	S138 and camera	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/04/24 15:24

EAST Search History (Interference)

< This search history is empty>

5/ 3/ 2019 2:45:17 PM
C:\ Users\ wwood\ Documents\ EAST\ Workspaces\ 15 808201.wsp

Applicant:	GORCHETCHNIKOV et. Al.
Title:	GRAPHIC PROECESSOR BASED ACCELERATOR SYSTEM AND METHOD
Appl. No.:	15/808,201
Filing Date:	11/09/2017
Examiner:	WILLIAM H. WOOD
Art Unit:	3992
Confirmation Number:	1077

Proposed Interview Agenda

Examiner Wood,

Thank you again for granting an interview on February 8 at 11 am ET. Here is a list of suggested topics for our discussion. If you have any questions about them before Friday, please feel free to call me at (617) 947-7280.

Sincerely,

Max Colice

Reg. No. 65,634

1 Information Disclosure Statement (IDS)

Would you please tell us why several of the references were not considered? If you need legible copies of the references that weren't considered, we would be happy to provide them.

2 Reissue Declaration

Would you please explain the defects in the declaration? We based it on a declaration accepted in another broadening reissue application, so we would appreciate your guidance on how it should be revised.

3 Recapture Rejections under 35 U.S.C. § 251

Here is a brief analysis of the 3-step recapture analysis in MPEP 1412.02(I) for each independent claim being added:

3.1 Claim 21

Step 1 – Yes, claim 21 is broader than claim 1 of the original patent (US 9,189,828; "the '828 patent")

Step 2 – Yes, the broader aspects relate to the "controller" in claim 1, which was relied upon to make claim 1 allowable over the prior art. (Applicant made other amendments to claim 1, but did not rely on those amendments to make claim 1 allowable over the prior art.)

Step 3 – Yes, claim 21 is narrower in a material respect than claim 1 because it recites "an output of a computation in the sequence of computations representing **an output of a neuron in an artificial neural network**." Per column 5, lines 20 – 25, of the '828 patent, an artificial neural network is a particular type of the sequence of computations claimed broadly in claim 1: "The elements of computation (computational elements) are characterized by a single output variable. Such computational elements often include one or more equations. Computational elements are same or similar within a large population and are computed in parallel. **An example of such a population is a layer of neurons in an artificial neural network (ANN), where all neurons are described by the same equation**." (Emphasis added.) This materially narrower limitation causes claim 21 to avoid recapture.

3.2 Claim 42

Step 1 – Yes, claim 42 is broader than corresponding claim 12 in the '828 patent.

Step 2 – No. Claim 12 in the '828 patent was not amended during prosecution, nor did Applicant argue about its scope. (Issued claim 12 had only double patenting rejections.) Therefore, Applicant did not surrender any subject matter related to claim 12 during prosecution.

Step 3 – Not applicable.

3.3 Claim 56

Step 1 – Yes, claim 56 is broader in at least some respects than claim 1 in the '828 patent.

Step 2 – Yes, although not every element was relied upon to distinguish prior art. *See Section 3.1 (Claim 21) above.*

Step 3 – Yes, claim 56 is narrower in a material respect than claim 1 because it recites "an output of a computation in the sequence of computations representing **an output of a neuron in an artificial neural network**." *See Section 3.1 (Claim 21) above.*

3.4 Claim 62

Step 1 – Yes, claim 62 is broader than claim 12.

Step 2 – No. Claim 12 in the '828 patent was not amended during prosecution, nor did Applicant argue about its scope. Issued claim 12 had only double patenting rejections.

Step 3 – Not applicable.

3.5 Claim 68

Recapture does not apply to claim 68 because claim 68 is direct to a different embodiment/ invention than the claims of the '828 patent. Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU accelerator and an architecture for the GPU accelerator, respectively. Claim 68, on the other hand, is directed to the embodiment of FIG. 3, which shows top-level control of the GPU accelerator system.

3.6 Claim 79

Recapture does not apply to claim 79 because claim 79 is direct to a different embodiment/ invention than the claims of the '828 patent. Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU

accelerator and an architecture for the GPU accelerator, respectively. Claim 79, on the other hand, is directed to the embodiment of FIG. 3, which shows top-level control of the GPU accelerator system.

3.7 Claim 90

Recapture does not apply to claim 90 because claim 90 is directed to a different embodiment/ invention than the claims of the '828 patent. Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU accelerator and an architecture for the GPU accelerator, respectively. Claim 90, on the other hand, is directed to the embodiment of FIG. 4, which shows the bottom level system and method performed during the computation on the GPU accelerator.

3.8 Claim 95

Recapture does not apply to claim 90 because claim 90 is directed to a different embodiment/ invention than the claims of the '828 patent. Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU accelerator and an architecture for the GPU accelerator, respectively. Claim 95, on the other hand, is directed to the embodiment of FIG. 4, which shows the bottom level system and method performed during the computation on the GPU accelerator.

3.9 Claim 101

Recapture does not apply to claim 90 because claim 90 is directed to a different embodiment/ invention than the claims of the '828 patent. Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU accelerator and an architecture for the GPU accelerator, respectively. Claim 101, on the other hand, is directed to the embodiment of FIG. 4, which shows the bottom level system and method performed during the computation on the GPU accelerator.

3.10 Claim 108

Recapture does not apply to claim 90 because claim 90 is directed to a different embodiment/ invention than the claims of the '828 patent. Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU accelerator and an architecture for the GPU accelerator, respectively. Claim 108, on the other

hand, is directed to the embodiment of FIG. 4, which shows the bottom level system and method performed during the computation on the GPU accelerator.

3.11 Claim 112

Recapture does not apply to claim 90 because claim 90 is directed to a different embodiment/ invention than the claims of the '828 patent. Claims 1 and 12 of the '828 patent are directed to the embodiments of FIGS. 1 and 2, which illustrate a computer system including the GPU accelerator and an architecture for the GPU accelerator, respectively. Claim 112, on the other hand, is directed to the embodiment of FIG. 4, which shows the bottom level system and method performed during the computation on the GPU accelerator.

3.12 Claim 117

Step 1 – Yes, claim 117 is broader than claim 1 because it omits the word "graphic" in "graphic processing unit."

Step 2 – No, because the word "graphic" was not added or argued during prosecution. It was in claim 1 as originally filed.

Step 3 – Not applicable.

3.13 Claim 128

Step 1 – Yes, claim 128 is broader than claim 12 because it omits the word "graphic" in "graphic processing unit."

Step 2 – No, because the word "graphic" was not added or argued during prosecution. It was in the original version of claim 12.

Step 3 – Not applicable.

4 Prior-Art Rejections

4.1 Claim 1

Buck does not teach "a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to transfer the at least a portion of the input data into the accelerator memory and to transfer at least a portion of the output data from the accelerator memory to the

main memory during performance of the sequence of computations by the at least one graphic processing unit" as in claim 1.

There is no controller element in Buck (see FIGS. 1 and 2).

4.2 Claim 12

Buck does not teach "**in parallel with performing the sequence of computations by the GPU in (A)**, transferring a second portion of the input data from the main memory to the accelerator via the bus; and **in parallel with performing the sequence of computations by the GPU in (A)**, transferring a second portion of the output data from the accelerator memory to the main memory via the bus." The training parameters are pre-loaded in the GPU and CPU in Buck. See column 10, lines 40–45, column 10, lines 60–65, and column 11, lines 25–30.

4.3 Claim 21

Buck does not teach "a controller, operably coupled to the GPU and the memory, to transfer the input data into the memory and to transfer at least a portion of the output data from the memory during performance of the sequence of computations by the GPU."

4.4 Claim 42

We propose overcoming the rejection by amending claim 42 to incorporate dependent claims 44 and 45.

4.5 Claim 56

We proposed overcoming the rejection by amending claim 56 to include dependent claim 61.

4.6 Claim 62

We propose amending claim 62 to include dependent claim 67

4.7 Claim 68

Buck does not teach "the user interaction stream controlling transfer of inputs to the first memory partition and the second memory partition; executing, by the processing unit, a computational stream, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations

representing the artificial neural network." The paragraphs of Buck cited in the Office Action state only:

- Buck, col. 9, lines 35–36: "The results of the computation are transferred from the machine learning GPU execution engine 200 to the CPU 120 for use in interactive use applications, such as speech recognition."
- Buck, col. 10, lines 41–43: "Data transfer from the CPU to the GPU is relatively expensive, so the CPU pre-loads as much as possible onto the GPU before entering a training loop."
- Buck, col. 8, lines 39–44: "The machine learning GPU execution engine and method disclosed herein alleviates the computational constraints of CPUs by porting some of the CPU processing to the GPU. More specifically, the present invention ports the machine learning architecture, traditionally handled by the CPU, to the GPU."

Buck is silent on a "user interaction stream" for controlling transfer of outputs or a "computational stream controlling data exchange between the user interaction stream and the computational stream."

4.8 Claim 79

Buck does not disclose or suggest "at least one central processing unit (CPU), operably coupled to the at least one memory, to execute a user interaction stream, the user interaction stream controlling transfer of inputs to the at least one memory; and a processing unit, operably coupled to the at least one memory and the at least one CPU, to execute a computational stream, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network."

4.9 Claim 90

Buck does not disclose or suggest "**in parallel with performing the sequence of computations** by the processing unit, transferring a second portion of the input data from a first memory partition to a second memory partition; and **in parallel with performing the sequence of computations** by the processing unit, transferring a second portion of the output data from the second memory partition to the first memory partition." *See Section 4.2 (Claim 12).*

4.10 Claim 95

Buck does not disclose or suggest "wherein the first memory partition is configured to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit, and wherein the second memory partition is configured to transfer a second portion of the output data to the first memory partition in parallel with performance the sequence of computations by the processing unit." *See Section 4.1 (Claim 1).*

4.11 Claim 101

We propose overcoming the rejection by amending claim 101 to recite "wherein swapping the first pointer and the second pointer includes synchronizing a user interaction stream and a computational stream."

4.12 Claim 108

Burton does not disclose or suggest "controlling performance of the at least one computation by the processing unit; controlling transfer of the output data from the processing unit to the memory during the first computational cycle; and controlling transfer of input data to the memory for use by the processing unit during a second computational cycle of the numerical simulation."

4.13 Claim 112

Burton does not teach a "a controller, operably coupled to the processing unit and the memory, to store the first output data in a second partition of the memory, the second partition referenced by a second pointer, and to swap the first pointer with the second pointer at the end of the first computational cycle, such that the first output data becomes an input for a second computational cycle of the numerical simulation."

4.14 Claim 117

Neither Buck nor Burton teaches "**a controller, operably coupled to the at least one processing unit and the accelerator memory**, to transfer the at least a portion of the input data into the accelerator memory and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one graphic processing unit."

4.15 Claim 128

Buck does not teach "(B) **in parallel with performing the sequence of computations by the processing unit in (A),** transferring a second portion of the input data from the main memory to the accelerator via the bus; and (C) **in parallel with performing the sequence of computations by the processing unit in (A),** transferring a second portion of the output data from the accelerator memory to the main memory via the bus."

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchechnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT/RESPONSE TO OFFICE ACTION

In response to the Office Action dated May 13, 2019, to which the deadline for responding is September 13, 2019 with a one-month extension of time, Applicant submits the following Amendments and Remarks, and respectfully requests reconsideration of the application in view thereof.

Amendments to the Claims are reflected in the listing of the claims which begins on page 2 of this paper.

Remarks begin on page 26 of this paper.

IN THE CLAIMS:

Set forth below in ascending order, with status identifiers, is a complete listing of all claims currently under examination. Changes to any amended claims are indicated by [[double brackets]], ~~strikethrough~~, and/or underlining. This listing also reflects any cancellation and/or addition of claims.

1. (Thrice Amended) A computer system, comprising:

a central processing unit to receive input data;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the [first] main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations yielding intermediate results; and

accelerator memory, operably coupled to the [graphic] at least one graphics processing unit, to store the results of the [plurality of sequential] sequence of computations; and

a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to initialize textures and shaders in the accelerator memory for performing the sequence of computations, to control performance of the sequence of computations by the at least one graphics processing unit, to transfer the at least a portion of the input data into the accelerator memory, and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one [graphic] graphics processing unit.

2. The computer system of claim 1, wherein the central processing unit is configured to receive the input data in response to a user interaction.

3. The computer system of claim 1, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one graphics processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

4. The computer system of claim 1, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

5. The computer system of claim 1, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

6. The computer system of claim 1, wherein accelerator memory comprises:

a first memory bank to store parameters common to all of the computations in the sequence of computations; and

a second memory bank to store data specific to at least one computation in the sequence of computations.

7. The computer system of claim 1, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

8. The computer system of claim 1, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one graphics processing unit has begun to perform another sequence of computations.

9. The computer system of claim 8, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one graphics processing unit.

10. The computer system of claim 1, wherein the controller is configured to control execution of the sequence of computations by the at least one graphics processing unit.

11. (Amended) The computer system of claim 1, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central [processor] <u>processing</u> unit, to acquire the input data in real time.

12. (Amended) A method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory, the method comprising:

(A) performing, by the GPU, the sequence of computations on a first portion of [the] input data so as to generate a first portion of [the] output data, intermediate computations in the sequence of computations yielding intermediate results, wherein performing the sequence of computations on the first portion of the input data comprises (i) assigning an output variable to a first texture and a second texture, the output variable being included in a first computational element of a plurality of computational elements, the plurality of computational elements representing the sequence of computations and (ii) accumulating a first value for the output variable in the first texture during a first time step;

(B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; [and]

(C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus; and

(D) performing, by the GPU, the sequence of computations on the second portion of the input data, wherein performing the sequence of computations on the second portion of the input data comprises (i) accumulating a second value for the output variable in the second texture during a second time step and (ii) making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements during the second time step.

13. The method of claim 12, further comprising:

storing the input data in the main memory in response to a user interaction.

14. The method of claim 12, further comprising:

receiving the input data at a first rate; and

wherein (A) comprises performing the sequence of computations at a second rate different than the first rate.

15. The method of claim 12, wherein (A) comprises:

generating an output representative of an output of at least one neuron in an artificial neural network.

16. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory without transferring any of the intermediate results of the plurality of sequential computations from the accelerator memory to the main memory so as to reduce data transfer via the bus.

17. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory after the GPU has begun to perform another sequence of computations.

18. The method of claim 17, wherein (C) further comprises:

initiating transfer of the second portion of the output data in parallel with performance of at least one computation in the other sequence of computations.

19. The method of claim 12, further comprising:

acquiring the input data in real time with at least one of a video camera, a microphone, or a cell recording electrode operably coupled to the CPU.

20. The method of claim 12, further comprising:

storing parameters common to all of the computations in the sequence of computations in a first memory bank in the accelerator memory; and

storing data specific to at least one computation in the sequence of computations in a second memory bank in the accelerator memory.

21. (New, Twice Amended) An accelerator system comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data, the sequence of computations representing an artificial neural network, computations in the sequence of computations representing respective layers of neurons in an artificial neural network;

a memory, operably coupled to the GPU, to store the input data and the output data; and

a controller, operably coupled to the GPU and the memory, to transfer the input data into the memory and to transfer at least a portion of the output data from the memory during performance of the sequence of computations by the GPU,

wherein each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture holding a first value computed by the neuron during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture holding a second value computed by the neuron during the current time step of the sequence of computations.

22. (New) The accelerator system of claim 21, wherein the GPU is operably coupled to a central processing unit (CPU) of a computer system.

23. (New) The accelerator system of claim 21, wherein the GPU is mounted on an expansion card operably coupled to a motherboard of a computer system.

24. (New) The accelerator system of claim 21, further comprising:

a central processing unit (CPU), a system memory, and a non-volatile memory installed on a motherboard.

25. (New) The accelerator system of claim 24, wherein the GPU is operably coupled to the CPU via a system bus on the motherboard.

26. (New) The accelerator system of claim 25, wherein the system bus has a bandwidth of 3.2 GB/s to 15 GB/s.

27. (New) The accelerator system of claim 25, wherein the controller is configured to transfer the input data and the at least a portion of the output data via the system bus.

28. (New) The accelerator system of claim 25, further comprising:

a video card, operably coupled to the motherboard via a local bus, to transfer graphical information to a display.

29. (New) The accelerator system of claim 28, wherein the local bus has a lower bandwidth than the system bus.

30. (New) The accelerator system of claim 21, wherein the memory comprises:

a first memory partition to store parameters common to every computation in the sequence of computations.

31. (New) The accelerator system of claim 30, wherein the memory further comprises:

a second memory partition to store data specific to the computation in the sequence of computations.

32. (New) The accelerator system of claim 31, wherein the second memory partition comprises:

a first partition to hold external input data patterns;

a second partition to hold representations of internal variables;

a third partition to hold an input of the computation in the sequence of computations; and

a fourth partition to hold the output of the computation in the sequence of computations.

33. (New) The accelerator system of claim 31, wherein the second memory partition is configured to accumulate the output of the computation in the sequence of computations during execution, by the GPU, of computational elements of the computation in the sequence of computations.

34. (New) The accelerator system of claim 31, wherein the second memory partition is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to computational elements executed by the GPU during the computation in the sequence of computations.

35. (New) The accelerator system of claim 21, wherein the computation in the sequence of computations performed by the GPU comprises a plurality of computational elements.

36. (New) The accelerator system of claim 35, wherein the plurality of computational elements represents a layer of neurons in an artificial neural network.

37. (New) The accelerator system of claim 36, wherein all neurons in the layer of neurons are described by the same equation.

38. (New) The accelerator system of claim 21, wherein the memory is configured to store intermediate results of the sequence of computations and the controller is configured to transfer the output data from the memory to a main memory without transferring the intermediate results.

39. (New) The accelerator system of claim 21, wherein the controller is configured to initiate transfer of at least a portion of the input data and to transfer the at least a portion of the output data in parallel.

40. (New) The accelerator system of claim 21, wherein the input data is acquired by at least one of a video camera, a microphone, or a cell recording electrode.

41. (New) The accelerator system of claim 21, wherein the input data is loaded from disk.

42. (New, Twice Amended) A method of performing a sequence of computations representing an artificial neural network, the method comprising:

receiving, at a central processing unit (CPU), first input data acquired from an external system;

initializing, by a controller operably coupled to a graphics processing unit (GPU), textures and shaders in a memory operably coupled to the GPU;

transferring the first input data received by the CPU to the memory operably coupled to the GPU;

performing, by the graphics processing unit (GPU), a first computation in the sequence of computations on the first input data based on the textures and shaders to generate first output data, computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of a computation in the sequence of computations representing an output of a neuron in the artificial neural network;

storing, in the memory operably coupled to the GPU, the first input data and the first output data; and

transferring second input data into the memory during performance of the sequence of computations by the GPU.

43. (New, Amended) The method of claim 42, wherein transferring the second input data comprises transferring the second input data via a bus operably coupled to the CPU.

44. (New) The method of claim 42, further comprising:

transferring the first output data from the memory to another memory during a second computation in the sequence of computations.

45. (New) The method of claim 44, further comprising:

storing intermediate results of the sequence of computations in the memory, and

wherein transferring the first output data from the memory to the other memory occurs without transferring the intermediate results of the sequence of computations.

46. (New) The method of claim 44, wherein transferring the second input data and transferring the first output data occurs in parallel.

47. (New) The method of claim 42, further comprising:

storing, in a first memory partition of the memory, parameters common to all of the computations in the sequence of computations.

48. (New) The method of claim 47, further comprising:

storing, in a second memory partition of the memory, data specific to the first computation in the sequence of computations.

49. (New) The method of claim 48, further comprising:

storing, in the second memory partition, external input data patterns, representations of internal variables, an input of the computation in the sequence of computations, and the output of the computation in the sequence of computations.

50. (New) The method of claim 42, wherein storing the first output data comprises:

accumulating, in the memory, outputs of computational elements executed by the GPU in performing the first computation in the sequence of computations.

51. (New) The method of claim 42, further comprising:

storing, in the memory, an output of a previous computation in the sequence of computations; and

accessing, by the GPU, the output of the previous computation during performance of the computation in the sequence of computations.

52. (New) The method of claim 42, wherein performing the first computation comprises executing a plurality of computational elements representing a layer of neurons in an artificial neural network.

53. (New) The method of claim 52, wherein all neurons in the layer of neurons are described by the same equation.

54. (New, Amended) The method of claim 42, further comprising:

acquiring the second input data with at least one of a video camera, a microphone, or a cell recording electrode.

55. (New, Twice Amended) The method of claim 42, further comprising:

loading the second input data from disk.

56. (New, Twice Amended) An accelerator system comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network;

a memory, operably coupled to the GPU, to store parameters common to all of the computations in the sequence of computations and data specific to the computation in the sequence of computations; and

a controller, operably coupled to the GPU and the memory, to control performance of the sequence of computations and data exchange by the GPU, to transfer at least a portion of the input data into the memory in parallel with performance of the sequence of computations by the GPU, and to transfer at least a portion of the output data from the memory in parallel with performance of the sequence of computations by the GPU.

57. (New) The accelerator system of claim 56, wherein all neurons in the layer of neurons are described by the same equation.

58. (New) The accelerator system of claim 56, wherein the memory is configured to hold:

external input data patterns;

representations of internal variables;

an input to the computation in the sequence of computations; and

the output of the computation in the sequence of computations.

59. (New) The accelerator system of claim 56, wherein the memory is configured to accumulate the output of the computation in the sequence of computations during execution of the plurality of computational elements by the GPU.

60. (New) The accelerator system of claim 56, wherein the memory is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to the plurality of computational elements during the computation in the sequence of computations.

61. Cancel claim 61.

62. (New, Twice Amended) A method comprising:

receiving, at a central processing unit (CPU), input data acquired from an external system;

initializing, by a controller, graphics processing unit (GPU) programs for performing a sequence of computations on the input data received by the CPU;

performing, with the GPU, the sequence of computations on the input data to generate output data based on the textures and shaders, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network;

controlling, by the controller, data exchange between the CPU and the GPU upon initialization of the GPU programs;

storing, in the memory operably coupled to the GPU, parameters common to all of the computations in the sequence of computations;

storing, in the memory, data specific to the computation in the sequence of computations; and

transferring at least a portion of the input data into the memory and at least a portion of the output data from the memory during performance of the sequence of computations by the GPU.

63. (New) The method of claim 62, wherein all neurons in the layer of neurons are described by the same equation.

64. (New) The method of claim 62, further comprising:

storing, in the memory, external input data patterns;

storing, in the memory, representations of internal variables; and

storing, in the memory, an input of the computation in the sequence of computations.

65. (New) The method of claim 62, further comprising:

accumulating, in the memory, the output of the computation in the sequence of computations during execution of the plurality of computational elements by the GPU.

66. (New) The method of claim 62, further comprising:

storing, in the memory, an output of a previous computation in the sequence of computations; and

accessing, by the GPU, the output of the previous computation during the computation in the sequence of computations.

67. Cancel claim 67

68. (New, Amended) A method of executing computations representing an artificial neural network on a computer system comprising at least one central processing unit (CPU), a processing unit, a first memory partition, and a second memory partition, the method comprising:

executing, by the at least one CPU, a user interaction stream, the user interaction stream controlling transfer of inputs to the artificial neural network to the first memory partition and the second memory partition; and

executing, by the processing unit, a computational stream, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network.

69. (New) The method of claim 68, wherein the user interaction stream controls data exchange between the user interaction stream and the computational stream outside of execution of the computations representing the artificial neural network.

70. (New) The method of claim 68, wherein executing the user interaction stream comprises:

controlling setting and editing of computational elements of the computations representing the artificial neural network.

71. (New) The method of claim 68, wherein executing the user interaction stream comprises:

controlling setting and editing of parameters of the computations representing the artificial neural network.

72. (New) The method of claim 68, wherein executing the user interaction stream comprises:

controlling setting and editing of parameters of the inputs to the artificial neural network.

73. (New) The method of claim 68, wherein executing the user interaction stream comprises:

specifying an output to be saved to disk and/or displayed on a screen.

74. (New) The method of claim 68, wherein executing the user interaction stream comprises:

parsing elements to be used in the computations representing the artificial neural network.

75. (New, Amended) The method of claim 74, wherein the processing unit comprises a graphics processing unit (GPU) and executing the user interaction stream further comprises:

converting the elements into GPU programs.

76. (New) The method of claim 75, wherein executing the user interaction stream comprises:

compiling the GPU programs.

77. (New) The method of claim 76, wherein executing the user interaction stream comprises:

transferring the GPU programs to the second memory partition.

78. (New) The method of claim 68, further comprising:

executing, by the at least one CPU, a data output stream, the data output stream controlling transfer of outputs of the computations representing the artificial neural network to disk.

79. (New, Amended) A system for executing computations representing an artificial neural network, the system comprising:

a first memory partition;

a second memory partition;

at least one central processing unit (CPU), operably coupled to the at least one memory, to execute a user interaction stream, the user interaction stream controlling transfer of inputs to the artificial neural network to the first memory partition and the second memory partition; and

a processing unit, operably coupled to the at least one memory and the at least one CPU, to execute a computational stream, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network.

80. (New) The system of claim 79, wherein the user interaction stream controls data exchange between the user interaction stream and the computational stream outside of execution of the computations representing the artificial neural network.

81. (New) The system of claim 79, wherein the user interaction stream controls setting and editing of computational elements of the computations representing the artificial neural network.

82. (New) The system of claim 79, wherein the user interaction stream controls setting and editing of parameters of the computations representing the artificial neural network.

83. (New) The system of claim 79, wherein the user interaction stream controls setting and editing of parameters of the inputs to the artificial neural network.

84. (New) The system of claim 79, wherein the user interaction stream specifies an output to be saved to disk and/or displayed on a screen.

85. (New) The system of claim 79, wherein the user interaction stream parses elements to be used in the computations representing the artificial neural network.

86. (New) The system of claim 85, wherein the processing unit comprises a graphics processing unit (GPU) and the user interaction stream converts the elements into GPU programs.

87. (New) The system of claim 86, wherein the user interaction stream compiles the GPU programs.

88. (New, Amended) The system of claim 87, wherein the user interaction stream transfers the GPU programs to at least one of the first memory partition or the second memory partition.

89. (New) The system of claim 68, wherein the at least one CPU is further configured to execute a data output stream, the data output stream controlling transfer of outputs of the computations representing the artificial neural network to disk.

90. (New, Twice Amended) A method of performing a sequence of computations with a processing unit, the method comprising:

receiving input data;

performing, by the processing unit, the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network, wherein performing the sequence of computations on the first portion of the input data comprises accumulating a first value for an output variable in a first texture by a first computational element;

in parallel with performing the sequence of computations on the first portion of the input data by the processing unit, transferring a second portion of the input data from a first memory partition to a second memory partition;

in parallel with performing the sequence of computations on the first portion of the input data by the processing unit, transferring a second portion of the output data from the second memory partition to the first memory partition; and

performing, by the processing unit, the sequence of computations on the second portion of the input data so as to generate a third portion of output data, wherein performing the sequence of computations on the second portion of the input data comprises accumulating a second value for the output variable in the first texture by the first computational element and accessing the first value of the output variable in the first texture by a second computational element.

91. (New) The method of claim 90, wherein receiving the input data comprises loading the input data from disk.

92. (New) The method of claim 90, wherein receiving the input data comprises receiving the input data from at least one of a video camera, a microphone, or a cell recording electrode.

93. (New) The method of claim 90, wherein receiving the input data occurs at a first rate and performing the sequence of computations occurs at a second rate different than the first rate.

94. (New) The method of claim 90, wherein transferring the second portion of the output data comprises transferring the second portion of the output data from the second memory partition to the first memory partition without transferring any of the intermediate results from the second memory partition to the first memory partition.

95. (New, Twice Amended) A system for performing a sequence of computations, the system comprising:

a first memory partition;

a second memory partition operably coupled to the first memory partition; and

a processing unit, operably coupled to the first memory partition and the second memory partition, to perform the sequence of computations on a first portion of input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network,

wherein the first memory partition is configured to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit,

wherein the second memory partition is configured to transfer a second portion of the output data to the first memory partition in parallel with performance the sequence of computations by the processing unit, and

wherein the sequence of computations represents an artificial neural network, each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture holds a first value of the output variable computed during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture accumulates a second value of the output variable computed during the current time step.

96. (New) The system of claim 95, wherein the first memory partition and the second memory partition are logical partitions.

97. (New) The system of claim 95, wherein the processing unit is comprises a graphics processing unit (GPU).

98. (New) The system of claim 95, wherein the processing unit is configured to receive the input data at a first rate and to perform the sequence of computations at a second rate is different than the first rate.

99. (New) The system of claim 95, wherein the second memory partition is configured to transfer the second portion of the output data to the first memory partition without transferring any of the intermediate results to the first memory partition.

100. (New) The system of claim 95, further comprising

at least one of a video camera, a microphone, or a cell recording electrode to generate the input data.

101. (New, Twice Amended) A method for performing a numerical simulation representing an artificial neural network over a plurality of computational cycles, the method comprising:

storing, in a first partition of a memory, first output data during the second computational cycle, the first output data generated for a first variable during the first computational cycle and being referenced by a first pointer, the first output data representing an output of a first neuron in the artificial neural network;

performing, by a processing unit operably coupled to the memory, at least one calculation on the first output data during the second computational cycle, the at least one calculation yielding second output data representing an output of the first neuron in the artificial neural network;

making the first output data in the first partition of the memory accessible to other neurons in the artificial neural network during the second computational cycle;

storing, in a second partition of the memory, the second output data, the second output data being referenced by a second pointer; and

swapping the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles.

102. (New) The method of claim 101, further comprising:

transferring input data into the memory from another memory coupled to a central processing unit (CPU) before the first computational cycle.

103. (New) The method of claim 101, further comprising:

providing, to a central processing unit (CPU), first output data generated during the first computational cycle after a conclusion of the first computational cycle.

104. (New) The method of claim 101, further comprising:

transferring the first output data from the memory to another memory during the second computational cycle.

105. (New) The method of claim 101, further comprising:

displaying, with a video system, an indication of the numerical simulation.

106. (New) The method of claim 101, further comprising:

dictating, to the processing unit, an order of execution of instructions.

107. (New) The method of claim 101, wherein performing the at least one calculation comprises receiving, by the processing unit, at least the portion of the input data and a series of instructions and executing, by the processing unit, the instructions on the at least the portion of the input data.

108. (New, Twice Amended) A method of performing a numerical simulation representing an artificial neural network, the method comprising:

performing, by a processing unit, at least one computation during a first computational cycle of the numerical simulation, the first computational cycle representing a first layer of neurons in the artificial neural network;

storing, by a memory operably coupled to the processing unit, data used to perform the at least one computation;

controlling performance of the at least one computation by the processing unit;

controlling transfer of the output data from the processing unit to the memory during the first computational cycle;

while performing the at least one computation, acquiring input data with a central processing unit and operably coupled to a main memory; and

controlling transfer of the input data to the memory from the main memory for use by the processing unit during a second computational cycle of the numerical simulation, the second computational cycle representing a second layer of neurons in the artificial neural network.

109. (New) The method of claim 108, further comprising:

dictating an order of execution of instructions to the processing unit.

110. (New, Amended) The method of claim 108, further comprising:

transferring output data from the memory to the main memory while the processing unit performs the at least one computation.

111. (New, Amended) The method of claim 108, further comprising:

performing successive computational cycles of the numerical simulation by feeding output data generated by the processing unit from a previous computational cycle and input data for a next computational cycle into the processing unit.

112. (New, Twice Amended) A system for executing an artificial neural network, the system comprising:

a central processing unit (CPU) to provide first input data;

a memory, operably coupled to the CPU, to store the first input data in a first partition, referenced by a first pointer, before computing a first layer of neurons of the artificial neural network;

a processing unit, operably coupled to the memory, to perform, during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data, the first output data representing an output of at least one neuron in the first layer of neurons; and

a controller, operably coupled to the processing unit and the memory, to:

store the first output data in a second partition of the memory, the second partition referenced by a second pointer, and to swap the first pointer with the second pointer at the end of the computation of the first layer of neurons, such that the first output data becomes an input for a second layer of neurons of the artificial neural network, and

dictate an order of execution of instructions to the processing unit to perform the computation of the first layer of neurons.

113. (New) The system of claim 112, wherein the processing unit comprises a graphics processing unit.

114. (New) The system of claim 112, wherein the controller is configured to send instructions for performing the at least one calculation to the processing unit.

115. (New, Amended) The system of claim 112, wherein the controller is further configured to transfer the first output data to another memory during computation of the second layer of neurons.

116. (New, Amended) The system of claim 112, wherein the memory further comprises:

a third partition to store internal variables; and

a fourth partition to store data used as input at a particular layer of neurons of the artificial neural network.

117. (New, Twice Amended) A computer system, comprising:

a central processing unit to receive input data acquired from an external system;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations yielding intermediate results; and

accelerator memory, operably coupled to the at least one processing unit, to store the results of the sequence of computations; and

a controller, operably coupled to the at least one processing unit and the accelerator memory, to control transfer of the at least a portion of the input data into the accelerator memory, to control transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one processing unit, and to control performance of the sequence of computations by the GPU.

118. (New) The computer system of claim 117, wherein the central processing unit is configured to receive the input data in response to a user interaction.

119. (New) The computer system of claim 117, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

120. (New) The computer system of claim 117, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

121. (New) The computer system of claim 117, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

122. (New) The computer system of claim 117, wherein accelerator memory comprises:

a first memory partition to store parameters common to all of the computations in the sequence of computations; and

a second memory partition to store data specific to at least one computation in the sequence of computations.

123. (New) The computer system of claim 117, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

124. (New) The computer system of claim 117, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one processing unit has begun to perform another sequence of computations.

125. (New) The computer system of claim 124, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one processing unit.

126. (New) The computer system of claim 117, wherein the controller is configured to control execution of the sequence of computations by the at least one processing unit.

127. (New, Amended) The computer system of claim 117, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time.

128. (New, Twice Amended) A method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a processing unit and an accelerator memory, the method comprising:

executing a first stream with the CPU, the first stream acquiring input data for the sequence of computations; and

executing a second stream with the processing unit, the second stream taking priority over the first stream during the sequence of computations and comprising:

(A) performing, by the processing unit, the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results;

(B) in parallel with performing the sequence of computations by the processing unit in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; and

(C) in parallel with performing the sequence of computations by the processing unit in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus.

129. (New) The method of claim 128, further comprising:

storing the input data in the main memory in response to a user interaction.

130. (New) The method of claim 128, further comprising:

receiving the input data at a first rate; and

wherein (A) comprises performing the sequence of computations at a second rate different than the first rate.

131. (New) The method of claim 128, wherein (A) comprises:

generating an output representative of an output of at least one neuron in an artificial neural network.

132. (New) The method of claim 128, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory without transferring any of the intermediate results of the plurality of sequential computations from the accelerator memory to the main memory so as to reduce data transfer via the bus.

133. (New) The method of claim 128, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory after the processing unit has begun to perform another sequence of computations.

134. (New) The method of claim 133, wherein (C) further comprises:

initiating transfer of the second portion of the output data in parallel with performance of at least one computation in the other sequence of computations.

135. (New) The method of claim 128, further comprising:

acquiring the input data in real time with at least one of a video camera, a microphone, or a cell recording electrode operably coupled to the processing unit.

136. (New) The method of claim 128, further comprising:

storing parameters common to all of the computations in the sequence of computations in a first memory partition in the accelerator memory; and

storing data specific to at least one computation in the sequence of computations in a second memory partition in the accelerator memory.

137. (New, Amended) A system for executing an artificial neural network, the system comprising:

a first processing unit to receive input data in real time from a sensor;

a memory, operably coupled to the first processing unit, to store a first portion of the input data;

a second processing unit, operably coupled to the memory, to generate first output data from the first portion of the input data, the first output data representing an output of at least one neuron in a first layer of neurons in the artificial neural network;

a controller, operably coupled to the second processing unit and the memory, to store a second portion of the input data in the memory while the second processing unit

generates the first output data and to allow a user to change the input data on the fly during the execution of the artificial neural network.

138. (New) The system of claim 137, wherein the sensor is a microphone.

139. (New) The system of claim 137, wherein the sensor is a video camera.

140. (New) The system of claim 137, wherein the second processing unit is further configured to generate second output data from the second portion of the input data, the second output data representing an output of at least one neuron in a second layer of neurons in the artificial neural network.

141. (New) The system of claim 140, further comprising:

a random access memory, operably coupled to the first processing unit, to receive the first output data from the memory while the second processing unit generates the second output data.

142. (New) The computer system of claim 1, wherein the controller is configured to inform the central processing unit that the sequence of computations is finished.

143. (New) The computer system of claim 1, wherein the controller is configured to reduce a processing load on the central processing unit.

144. (New) The computer system of claim 1, wherein the controller is configured to reduce interactions between the central processing unit and the accelerator.

145. (New) The method of claim 68, wherein the computational stream takes priority over the user interaction stream during execution of the computations representing the artificial neural network

146. (New) The method of claim 68, wherein the computational stream determines a timing of a transfer of input data from the CPU to the processing unit and a timing of a transfer of output data from the processing unit to the CPU during execution of the computations representing the artificial neural network.

147. (New) The method of claim 108, wherein acquiring the input data comprises acquiring the input data in real time with a video camera.

148. (New) The method of claim 128, wherein the sequence of computations represents an artificial neural network, each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, and wherein (A) further comprises, during a current time step of the sequence of computations:

storing a first value of the output variable computed during a previous time step of the sequence of computations in a first portion of the accelerator memory;

making the first value accessible to every neuron in the neural network during a current time step of the sequence of computations; and

accumulating a second value of the output variable computed during the current time step in a second portion of the accelerator memory.

REMARKS

In response to the Office Action dated May 13, 2019, in the above-identified application, Applicant respectfully requests reconsideration. All of the issues raised in the Office Action have been carefully considered and are addressed in this response. The application as now presented is believed to be in allowable condition.

1. Status of the Claims and Summary of Amendments

Upon entry of this response, claims 1–60, 62–66, and 68–148 will be pending in this application, of which claims 1, 12, 21, 42, 56, 62, 68, 79, 90, 95, 101, 108, 112, 117, 128, and 137 are independent claims.

Claim 1 is being amended to recite "a controller ... *to initialize textures and shaders in the accelerator memory for performing the sequence of computations, to control performance of the sequence of computations by the at least one graphics processing unit,* to transfer the at least a portion of the input data into the accelerator memory, and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one graphics processing unit." Support for this amendment appears at least in column 5, lines 1–18, and in column 8, lines 5–33 and lines 45–60 of U.S. Patent No. 9,189,828 ("the '828 patent").

Claim 12 is being amended to recite "performing the sequence of computations on the first portion of the input data comprises (i) assigning an output variable to a first texture and a second texture, the output variable being included in a first computational element of a plurality of computational elements, the plurality of computational elements representing the sequence of computations and (ii) accumulating a first value for the output variable in the first texture during a first time step" and "performing, by the GPU, the sequence of computations on the second portion of the input data, wherein performing the sequence of computations on the second portion of the input data comprises (i) accumulating a second value for the output variable in the second texture during a second time step and (ii) making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements, during the second time step." Support for this amendment appears at least in column 6, lines 3–17 of the '828 patent.

Claim 21 is being amended to recite "each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture

holding a first value computed by the neuron during a previous time step and *accessible to other neurons in the neural network during a current time step* and the second texture holding a second value computed by the neuron during the current time step." Support for these amendments appears at least in column 5, line 61, through column 6, line 17 of the '828 patent.

Claim 42 is being amended to recite "*initializing, by a controller operably coupled to a graphics processing unit (GPU), textures and shaders in a memory operably coupled to the GPU*; transferring the first input data received by the CPU to the memory operably coupled to the GPU; [and] performing, by the graphics processing unit (GPU), a first computation in the sequence of computations on the first input data *based on the textures and shaders*." Support for these amendments appears at least in column 5, line 61, through column 6, line 17 of the '828 patent.

Claim 56 is being amended to recite "a controller, operably coupled to the GPU and the memory, to control performance of the sequence of computations and data exchange by the GPU." Support for these amendments appears at least in column 5, lines 1–12, and column 8, lines 5–33 and lines 45–60 of the '828 patent.

Claim 62 is being amended to recite "initializing, by a controller, graphics processing unit (GPU) programs for performing a sequence of computations on the input data received by the CPU" and "controlling, by the controller, data exchange between the CPU and the GPU upon initialization of the GPU programs." Support for these amendments appears at least in column 5, lines 1–12, and column 8, lines 5–33 and lines 45–60 of the '828 patent.

Claim 90 is being amended to recite "each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture holds a first value of the output variable computed during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture accumulates a second value of the output variable computed during the current time step. Support for these amendments appears at least in column 5, line 61, through column 6, line 17 of the '828 patent.

Claim 95 is being amended to recite "the sequence of computations represents an artificial neural network, each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture holds a first value computed during a previous time step and accessible to other neurons in the neural network during a current time step, and the second texture holds a second value computed during the

current time step." Support for this amendment appears at least in column 5, line 61, through column 6, line 17 of the '828 patent.

Claim 101 is being amended for clarity and to recite, among other things, "storing, in a first partition of a memory, first output data during the second computational cycle, the first output data generated for a first variable during the first computational cycle and being referenced by a first pointer, *the first output data representing an output of a first neuron in the artificial neural network*" and "accessing, by a second neuron in the artificial neural network during the second computational cycle, the first output data in the first partition of the memory." Support for this amendment appears at least in column 5, line 61, through column 6, line 17 of the '828 patent.

Claim 108 is being amended to recite "while performing the at least one computation, acquiring input data with a central processing unit and operably coupled to a main memory." Support for this amendment appears at least in column 9, lines 8–18 of the '828 patent.

Claim 112 is being amended to recite "a controller, operably coupled to the processing unit and the memory, to ... dictate an order of execution of instructions to the processing unit to perform the computation of the first layer of neurons." Support for this amendment appears at least in column 2, lines 47–52 of the '828 patent. Support for this amendment appears at least in column 5, lines 1–18, and in column 8, lines 5–33 and lines 45–60 of the '828 patent.

Claim 117 is being amended to recite "a controller ... to control transfer of the at least a portion of the input data into the accelerator memory, to control transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one processing unit, and to control performance of the sequence of computations by the GPU." Support for this amendment appears at least in column 5, lines 1–18, and in column 8, lines 5–33 and lines 45–60 of the '828 patent.

Claim 128 is being amended to recite "executing a first stream with the CPU, the first stream acquiring input data for the sequence of computations," "executing a second stream with the processing unit," and that the second stream "tak[es] priority over the first stream during the sequence of computations." Support for these amendments appears at least in column 7, line 31

Claim 137 is being amended to read "a controller, operably coupled to the second processing unit and the memory, to store a second portion of the input data in the memory while the second processing unit generates the first output data, and *to allow a user to change the input*

data on the fly during the execution of the artificial neural network." Support for this amendment appears at least in column 9, lines 8–18 of the '828 patent.

Claim 142, which depends from independent claim 1 and recites that "the controller is configured to inform the central processing unit that the sequence of computations is finished," is being added. Support for claim 142 appears at least in column 9, lines 19–21 of the '828 patent.

Claim 143, which depends from independent claim 1 and recites that "the controller is configured to reduce a processing load on the central processing unit," is being added. Support for claim 143 appears at least in column 13, lines 35–42 of the '828 patent.

Claim 144, which depends from independent claim 1 and recites that "the controller is configured to reduce interactions between the central processing unit and the accelerator," is being added. Support for claim 144 appears at least in column 13, lines 35–42 of the '828 patent.

Claim 145, which depends from independent claim 68 and recites that "the computational stream takes priority over the user interaction stream during execution of the computations representing the artificial neural network," is being added. Support for claim 145 appears at least in column 7, lines 45–55 of the '828 patent.

Claim 146, which depends from independent claim 68 and recites that the computational stream controls timing of data transfers between the CPU and the processing unit during execution of the computations representing the artificial neural network, is being added. Support for claim 146 appears at least in column 13, lines 1–10 of the '828 patent.

Claim 147, which depends from independent claim 108 and recites that "acquiring the input data comprises acquiring the input data in real time with a video camera," is being added. Support for claim 144 appears at least in column 9, lines 8–18 of the '828 patent.

Claim 148, which depends from claim 128, is being added. Support for claim 148 appears at least in column 5, line 61, through column 6, line 17 of the '828 patent.

2. Information Disclosure Statement

According to page 3 of the Office Action, Item 1[1] cited in the information disclosure statement (IDS) filed on March 20, 2019, could not be found. Applicant apologizes for this mistake and is submitting a copy of Item 1 in the attached IDS. Applicant notes that Item 1 is the abstract of a talk given by one of the inventors after the priority date of the '828 patent.

[1] "Brain-inspired computing: Invited keynote address," Bionetics 2010, Boston MA, USA. 1 page.

Page 3 of the Office Action states also that a document entitled "From Spikes to Inter-Areal Synchrony: How attentive matching and resonance control learning and information processing by laminar thalamocortical circuits" was included with the IDS filed March 20, 2019. This one-page document is cumulative of a 2008 article submitted and considered as Item 96 on the IDS filed December 8, 2017.

3. Objections to Claims

On page 3 of the Office Action, claim 12 was objected to because the amendment submitted on March 20, 2019 ("the March 2019 amendment") did not contain a parenthetical expression as required by 37 C.F.R. 1.173(b)(2) or 1.173(c). Applicant is amending claim 12 to include the parenthetical expression and address this objection.

Claims 21–54, 62–66, 101–127, and 137–141 were objected to because the March 2019 amendment did not provide support for the changes as required by 37 C.F.R. 1.173(c).

In particular, claims 21–41 were objected to because the March 2019 amendment allegedly does not show "computations represent layers."[2] Applicant respectfully disagrees. The March 2019 amendment states that support for "computations in the sequence of computations represent respective layers of neurons in an artificial neural network" can be found at least in "column 5, lines 22–26 and 61–63 and column 9, lines 3–7"[3] of the '828 patent. Column 5, lines 22–26 of the '828 patent are reproduced below:

> "The elements of computation (computational elements) are characterized by a single output variable. Such computational elements often include one or more equations. Computational elements are same or similar within a large population and are computed in parallel. An example of such a population is a layer of neurons in an artificial neural network (ANN), where all neurons are described by the same equation."

As described in the section of the '828 patent quoted above, computational elements can include computations (e.g., equations) that are the same or similar within a large population. An example of such a population is a layer of neurons in an artificial neural network. In other words, computations can represent a layer of neurons in an artificial neural network. Therefore, the March 2019 amendment identifies support as required by 37 C.F.R. 1.173(c).

[2] Office Action, page 3.
[3] March 2019 amendment, page 25.

Claims 42–54 were objected to because the March 2019 amendment did not contain support for the changes to independent claim 42. Independent claim 42 was amended in the March 2019 amendment to recite in part "computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of a computation in the sequence of computations representing an output of a neuron in the artificial neural network." Support for this change can be found at least in column 5, lines 22–26 and 61–63 and column 9, lines 3–7 of the '828 patent as explained immediately above.

Claims 42–54, 62–66, and 117–127 were objected to because the March 2019 amendment allegedly did not provide support for the claim language "external system." Independent claims 42, 62, and 117 recite in part "receiving, at a central processing unit (CPU), first input data acquired from an external system." Applicant submits that support for this limitation can be found at least in column 8, lines 6–12 (the "recording device" is an external system); column 9, lines 8–13 (giving examples of recording devices); and column 12, lines 54–67 (CPU receives external input) of the '828 patent.

Claims 101–116 and 137–141 were objected to because the March 2019 amendment allegedly did not contain support for claim language relating to "computational cycles to first and second layers of neurons."[4] The March 2019 amendment states that support for "a first computational cycle representing a first layer of neurons in the artificial neural network" and "a second computational cycle representing a second layer of neurons in the artificial neural network" can be found at least in "column 1, lines 38–41, column 5, lines 22-26 and column 9, lines 3–7"[5] of the '828 patent. Column 5, lines 22–26 are reproduced below and indicate that computational elements may represent layers of neurons in an artificial neural network:

> The elements of computation (computational elements) are characterized by a single output variable. Such computational elements often include one or more equations. Computational elements are same or similar within a large population and are computed in parallel. An example of such a population is a layer of neurons in an artificial neural network (ANN), where all neurons are described by the same equation.

And column 11, lines 61-65, and column 12, lines 9-17, describe the relationship between computations and computational cycles.

[4] Office Action, page 4.
[5] March 2019 amendment, page 26.

Additional support for the claimed computational cycles and layers of neurons appears in col. 9, line 25, through col. 11, line 20 of the '828 patent. This section of the '828 patent describes the implementation of the Synchronous Artificial Neuronal Network Distributed Runtime Algorithm (SANNDRA) with a GPU as in claims 101-116 and 137-141 and involves creating many neuron layers, each of which is executed during a corresponding computational cycle. Therefore, the '828 patent includes support for first and second computational cycles representing first and second layers of neurons, respectively.

Accordingly, withdrawal of the objections is respectfully requested.

4. Rejections Under 35 U.S.C. § 112

Claims 62–66 were rejected as indefinite. According to page 4 of the Office Action, the phrase "transferring the input data into the memory ... during performance of the sequence of computations by the GPU" in independent claim 62 is "being treated as stating, 'transferring at least a portion of the input data ...'." Page 4 of the Office Action states further that "[t]his interpretation is necessary in order for at least a portion of the input data to have already been transferred so that in can be performed upon by the GPU."

As indicated above, Applicant is amending independent claim 62 to match the Office's interpretation. Therefore, withdrawal of the indefiniteness rejection is respectfully requested.

5. Rejections Under 35 U.S.C. § 102(e)

Claims 1–7, 10–16, 19–40, 42, 43, 47–54, 56–60, 62–66, 68–72, 74–77, 79–83, 85–88, 90, 92–95, 97–100, 117–123, 126–132, 135–138, 140, and 141 (including independent claims 1, 12, 21, 42, 56, 62, 68, 79, 90, 95, 117, 128, and 137) were rejected as anticipated by U.S. Patent 7,219,085 to Buck. Applicant respectfully traverses these rejections to the extent they are maintained over the claims as amended herein. (For brevity, Applicant is addressing only the rejections of the independent claims below.)

Claim 1

Claim 1, which recites a CPU, main memory, GPU, accelerator memory, and controller, was rejected as anticipated by Buck. As indicated above, claim 1 is being amended to specify that the controller "initialize[s] textures and shaders in the accelerator memory for performing the sequence of computations" and "control[s] performance of the sequence of computations by

the [GPU]" in addition to transferring data between the main memory and the accelerator memory while the GPU performs the sequence of computations.

To anticipate amended claim 1, Buck would have to disclose the controller of amended claim 1 in as much detail as recited in amended claim 1. In this case, however, Buck doesn't disclose such a controller. At best, Buck discloses some of the functions performed by the controller of amended claim 1 are performed by the CPU or are simply performed without specifying what performs those functions. As explained below, performing these functions with a controller as in claim 1 instead of a CPU as in Buck leads to significant performance advantages.

To start, Buck states that the CPU controls performance of the sequence of computations by the GPU: "Once the training data (X') has been loaded on the GPU, *the CPU instructs the GPU to run the various shaders need 30 Testing Phase to process the machine learning technique* (box 540)."[6] Similarly, Buck states that the CPU transfers input data during the training phase: "The CPU then starts to loop on the training data and accumulates groups of training data by loading batches of training data (X,T) (box 515)."[7] In contrast, claim 1 recites that the controller of "control[s] performance of the sequence of computations by the [GPU]" and transfers data between the main memory and the accelerator memory while the GPU performs the sequence of computations.

Buck discloses other functions performed by the claimed controller without explaining which components perform those functions. For example, when Buck describes loading shaders, Buck doesn't say *what* loads the shaders. Instead, Buck just says that the shaders *are* loaded: "[T]he training phase of the machine learning GPU execution method begins by loading shaders programs (P), initial trainable parameters (W), such as weights, and other learning parameters (L) (box 500). This data is then transferred from the CPU to the GPU (arrow 505)."[8]

Similarly, Buck discloses transferring output data (training statistics) while the GPU performs a machine learning training phase but doesn't say which component is responsible for the transfer. Instead, Buck describes the transfer using the passive voice: "[A]s shown in FIG. 5, the CPU can collect training statistics and obtain training progress data (box 560). *The* [sic] *is achieved by having the GPU send the CPU* the updated trainable parameters (W) and the errors

[6] Buck, 11:29–31 (emphasis added).
[7] Buck, 10:51–53.
[8] Buck, 10:36–41.

(arrow 565)."[9] This makes it difficult, if not impossible, to determine which component controls the data transfer.

Thus, Buck does not provide the controller of amended claim 1.

Buck's use of the CPU to perform functions performed by the controller of claim 1 illustrates a fundamental difference between Buck's system and the system of claim 1: in Buck's system, the CPU controls computations by and interactions with the GPU, whereas in the claimed system, the controller controls computations by and interactions with the GPU. As explained in the '828 patent, using a controller instead of a CPU to load textures and shaders, control GPU computations, and transfer data reduces the processing load on the CPU, so the CPU "can be dedicated to other tasks."[10] Put differently, using a controller to perform these functions "frees the CPU 120 to execute other applications and allows the expansion card [the accelerator of amended claim 1] to run at its full capacity without being slowed down by extensive interactions with the CPU."[11] These are significant technical advantages to using the CPU to perform the same functions.

Claim 12

Claim 12 was also rejected as anticipated by Buck. As indicated above, claim 12 is being amended to recite "performing the sequence of computations on the first portion of the input data comprises (i) assigning an output variable to a first texture and a second texture, the output variable being included in a first computational element of a plurality of computational elements, the plurality of computational elements representing the sequence of computations and (ii) accumulating a first value for the output variable in the first texture during a first time step" and "performing, by the GPU, the sequence of computations on the second portion of the input data, wherein performing the sequence of computations on the second portion of the input data comprises (i) accumulating a second value for the output variable in the second texture during a second time step and (ii) making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements during the second time step."

[9] Buck, 11:61–65 (emphasis added).
[10] '828 patent, Abstract.
[11] '828 patent, 13:39–42.

Buck is silent on assigning an output variable associated with a first computational element to first and second textures. Buck is also silent on accumulating a first value of the output variable in the first texture during a first time step, then making that first value accessible to other computational elements during a second time step. And Buck is silent on accumulating a second value of the output variable in the second texture during the second time step. Therefore, Buck doesn't anticipate amended claim 12.

Moreover, amended claim 12 isn't obvious over Buck in view of US 2001/0010034 A1 to Burton, which was also cited in the Office Action and which discloses simulating pipelined hardware.[12] Although Burton discloses writing outputs to different memory areas on different clock cycles of the simulation,[13] Burton doesn't disclose or suggest making those memory areas accessible to other computational elements after they have been written. At best, Burton just says that the memory areas can be part of a common memory.[14] Therefore, Burton doesn't cure Buck's deficiencies with respect to amended claim 12.

Claim 21

Claim 21 was also rejected as anticipated by Buck. As indicated above, claim 21 is being amended to recite "each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, *the first texture holding a first value computed output by the neuron during a previous time step and accessible to other neurons in the neural network during a current time step* and the second texture holding a second value computed by the neuron during the current time step." Buck does not disclose assigning each neuron in an artificial neural network to first and second textures or allowing every neuron to access the first texture while computing a value for the second texture during the same time step. Therefore, Buck doesn't anticipate claim 21 as amended.

Claim 42

Claim 42 was rejected as anticipated by Buck. As indicating above, claim 42 is being amended to recite "*initializing, by a controller operably coupled to a graphics processing unit*

[12] Burton, Abstract.
[13] *See, e.g.,* Burton, paragraphs 13–16.
[14] Burton, paragraph 74 ("The first data storage area 2 and the second data storage area 4 are respective regions within the common memory 24. The local variable storage 6 is also provided within the common memory 24. The software models of the pipelined circuit elements 26 are also stored within the common memory 24 together with the overall controlling software for the simulation system as a whole. The swapping of the pointers to the memory areas 2, 4 can effectively be achieved by swapping memory address offset values or similar techniques."); *see also* Burton, paragraph 79.

(GPU), textures and shaders in a memory operably coupled to the GPU." As explained above with respect to claim 1, Buck doesn't disclose using a controller to initialize textures and shaders in a memory operably coupled to a GPU. Instead, Buck says that the textures and shaders are "loaded" and implies that the CPU handles this initialization.[15] Initializing textures and shaders with a CPU is not the same as or equivalent to initializing textures and shaders with a controller. Using a controller reduces the processing load on the CPU and allows the GPU to run at its full capacity.[16] Given these differences, claim 42 as amended is patentable over Buck.

Claim 56

Claim 56 was rejected as anticipated by Buck and is being amended to recite "a controller, operably coupled to the GPU and the memory, to control performance of the sequence of computations and data exchange by the GPU." As explained above with respect to amended claims 1 and 42, Buck does not disclose or suggest such a controller. Thus, amended claim 56 is patentable over Buck.

Claim 62

Claim 62 was rejected as anticipated by Buck and is being amended to recite "initializing, *by a controller*, graphics processing unit (GPU) programs for performing a sequence of computations on the input data received by the CPU" and "controlling, *by the controller*, data exchange between the CPU and the GPU upon initialization of the GPU programs." As explained above with respect to amended claims 1, 42, and 56, Buck does not disclose or suggest such a controller. Thus, amended claim 62 is patentable over Buck.

Claim 68

Claim 68, which recites a user interaction stream executed by a CPU and computation stream executed by a processing unit, was rejected as anticipated by Buck. To anticipate claim 68, Buck would have disclose each and every element of claim 68.[17] In this case, however, Buck doesn't disclose the computational stream of claim 68. Therefore, Buck doesn't anticipate claim 68.

---

[15] Buck, 10:36–41 ("[T]he training phase of the machine learning GPU execution method begins by loading shaders programs (P), initial trainable parameters (W), such as weights, and other learning parameters (L) (box 500). This data is then transferred from the CPU to the GPU (arrow 505).").
[16] '828 patent, 13:39–42.
[17] *See* MPEP 2131.

Claim 68 specifies that the computational stream is executed by a processing unit separate from the CPU and "control[s] data exchange between the user interaction stream and the computational stream during execution of the computations representing [an] artificial neural network." Page 30 of the Office Action identifies three sections of Buck as providing this computational stream: "executing, by the processing unit, a computational stream *(Buck: [1] column 2, lines 24-28; [2] column 8, lines 39-44)*, the computational stream controlling data exchange between the user interaction stream and the computational stream during execution of the computations representing the artificial neural network *(Buck: [3] figure 5, elements 220 and 545)*."[18] But these sections don't disclose or suggest a computational stream that controls data exchange as recited in claim 68.

The first cited section of Buck—col. 2, lines 24–28—says merely that GPUs are stream processors: "For general purpose computing, GPUs are essentially stream processors with limitations. Dealing with and circumventing these limitations requires a style of programming and processing that is neither obvious nor intuitive." This section doesn't mention any particular stream, nor does it disclose how the streams operate or what they do, much less that a processing unit other than a CPU executes a computational stream that controls data exchange between a user interaction stream and the computational stream during execution of an artificial neural network. In fact, the first cited section of Buck is silent on data exchanges.

The second cited section of Buck—col. 8, lines 39–44—just says that the GPU performs some of the machine learning that would otherwise be performed by the CPU: "The machine learning GPU execution engine and method disclosed herein alleviates the computational constraints of CPUs by porting some of the CPU processing to the GPU. More specifically, the present invention ports the machine learning architecture, traditionally handled by the CPU, to the GPU." Like the first cited section, the second cited section doesn't say anything about streams or controlling data exchanges between streams.

The third cited section of Buck— elements 220 and 545 in Fig. 5—fails to provide the pieces of the claimed computational stream missing from the first and second cited sections of Buck. Element 220 in Buck's Fig. 5 is a GPU. And element 545 in Buck's Fig. 5 is a learning module. Buck's Fig. 5 and the associated description don't disclose anything about computational streams or controlling data exchanges between streams while executing

[18] Italics original.

computations representing an artificial neural network as in claim 68. Therefore, these sections of Buck don't provide the "computational stream" step of claim 68.

The rest of Buck is also silent on a computational stream, executed by a processing unit other than a CPU, that controls data exchange between user interaction and computational stream during execution of an artificial neural network as in claim 68. Although Buck discloses transferring data between the CPU and a GPU while the GPU executes a neural network, Buck is silent on which component controls the data exchange.[19] Thus, Buck doesn't anticipate claim 68.

Claim 142 depends from claim 68 and is patentable over Buck for additional reasons as well. In particular, claim 142 recites that "the computational stream takes priority over the user interaction stream during execution of the computations representing the artificial neural network." As Buck doesn't disclose the computational stream of claim 68, Buck can't disclose that the computational stream takes priority over another stream as in claim 142.

Claim 79

Like claim 68, claim 79 was rejected as anticipated by Buck. But like claim 68, claim 79 recites "a processing unit, operably coupled to the at least one memory and the at least one CPU, to execute a computational stream, *the computational stream controlling data exchange between the user interaction stream and the computational stream*." Thus, like claim 68, claim 79 is patentable over Buck.

Claim 90

Like claim 21, claim 90 was rejected as anticipated by Buck. And like claim 21, claim 90 is being amended to recite "a first texture" that "holds a first value computed during a previous time step and accessible to other neurons in the neural network during a current time step" and "a second texture" that "holds a second value computed during the current time step." Thus, like amended claim 21, amended claim 90 is patentable over Buck.

Claim 95

Like claim 21, claim 95 was rejected as anticipated by Buck. And like claim 21, claim 95 is being amended to recite "a first texture" that "holds a first value computed during a previous time step and accessible to other neurons in the neural network during a current time step" and "a

[19] *See, e.g.*, Buck, 11:63–65 and 12:8–10 (using the passive voice to describe data transfers between the CPU and the GPU).

second texture" that "holds a second value computed during the current time step." Thus, like amended claim 21, amended claim 95 is patentable over Buck.

Claim 117

Claim 117 was rejected as anticipated by Buck and is being amended to recite "a controller ... to control transfer of the at least a portion of the input data into the accelerator memory, to control transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one processing unit, and to control performance of the sequence of computations by the GPU." Claim 117 as amended is therefore patentable over Buck for at least some of the reasons given above with respect to claim 1.

Claim 128

Claim 128 was rejected as anticipated by Buck and is being amended to recite "executing a first stream with the CPU, the first stream acquiring input data for the sequence of computations," "executing a second stream with the processing unit," and that the second stream "tak[es] priority over the first stream during the sequence of computations." As explained above with respect to claim 68, Buck is silent on a second (computational) stream executed by a processing unit that involves transferring data while performing a sequence of computations. Buck is also silent on one stream taking priority over another stream. Therefore, Buck doesn't anticipate amended claim 128.

Claim 137

Claim 137, which recites a first processing unit that receives input data in real time from a sensor, was rejected as anticipated by Buck. Page 33 of the Office Action implies that Buck's disclosure of different types of input devices, including a microphone,[20] provided the "sensor" of claim 137. And as best understood, page 33 of the Office Action equates Buck's disclosure of interactive use applications in col. 9, lines 33–36 to provide "receiv[ing] input data in real time" as in claim 137.

Applicant respectfully submits that this interpretation of Buck's interactive use applications is incorrect. For completeness, here is col. 9, lines 33–36 of Buck: "The results of

[20] Buck, 7:51 – 8:2;

the computation are transferred from the machine learning GPU execution engine 200 to the CPU 120 for use in interactive use applications, such as speech recognition." This section of Buck refers to "the results of the computation," which a reader of ordinary skill in art would understand to be the weights for the trained neural network or machine learning architecture. These weights are used in the trained and tested neural network, which is executed by the CPU after training and testing on the GPU are complete, not in the training or testing carried out by the GPU. Buck is silent on whether the interactive use application provides any input data, in real-time, for training or testing the GPU. Therefore, Buck doesn't anticipate claim 137.

Nevertheless, solely in an effort to advance prosecution, Applicant is amending claim 137 to recite a controller that can "allow a user to change the input data on the fly during the execution of the artificial neural network." Although Buck discloses providing input data to the GPU for training and testing a neural network,[21] Buck doesn't disclose that a user can change this input data while the GPU is executing the neural network. Thus, amended claim 137 is patentable for additional reasons as well.

6. Rejections Under 35 U.S.C. § 103(a) Over Buck in view of Burton

Claims 101–104 and 106–116 were rejected as unpatentable over Buck in view of U.S. Patent Publication No. 2001/0010034 to Burton. Applicant respectfully traverses these rejections to the extent they are maintained over the claims as amended herein. (For brevity, Applicant is addressing only the rejections of the independent claims below.)

Claim 101

Claim 101 recites storing first output data from a first neuron in a neural network in a first memory partition during a first computational cycle, storing second output data from the first neuron in a second memory partition during a second computational cycle, then swapping pointers referencing the first and second memory partitions before a third computational cycle. According to pages 40–42 of the Office Action, Burton provides every element of claim 101 except the artificial neural network, which is provided by Buck. The Office Action states further that using Burton's pipelining techniques with Buck's artificial neural network would have been nothing more than "a substitution and application of one known technique and element for another yielding a predictable result."[22]

[21] *See, e.g.*, Buck, 11:59–12:6.
[22] Office Action, pages 41–42.

Like claim 12, claim 101 is being amended to recite "making the first output data in the first partition of the memory accessible to other neurons in the artificial neural network during the second computational cycle." As explained above with respect to claim 12, Burton discloses writing outputs to different memory areas on different clock cycles of the simulation,[23] but is silent on making those memory areas accessible to other computational elements after they have been written. At best, Burton just says that the memory areas can be part of a common memory.[24]

Buck's artificial neural network doesn't cure Burton's deficiencies with respect to amended claim 101. Therefore, amended claim 101 is patentable over the asserted combination of Burton and Buck.

Claim 108

Claim 108 recites a method of performing a numerical simulation representing an artificial neural network. According to pages 43 and 44 of the Office Action, Burton provides every element of claim 108 except the artificial neural network, which is provided by Buck. The Office Action states further that using Burton's pipelining techniques with Buck's artificial neural network would have been nothing more than "a substitution and application of one known technique and element for another yielding a predictable result."[25]

As indicated above, Applicant is amending claim 108 to recite, "while performing the at least one computation [with the processing unit during a first computational cycle of the numerical simulation], acquiring input data with a central processing unit and operably coupled to a main memory." Claim 108 recites further controlling the transfer of this input data for use by the processing unit during a second computational cycle of the numerical simulation. Burton and Buck are each silent on acquiring input data while performing computations, much less transferring that input data to a processing unit (e.g., a GPU) for use in the computations. Accordingly, no combination of Burton and Buck yields the method of amended claim 108.

---

[23] *See, e.g.,* Burton, paragraphs 13–16.
[24] Burton, paragraph 74 ("The first data storage area 2 and the second data storage area 4 are respective regions within the common memory 24. The local variable storage 6 is also provided within the common memory 24. The software models of the pipelined circuit elements 26 are also stored within the common memory 24 together with the overall controlling software for the simulation system as a whole. The swapping of the pointers to the memory areas 2, 4 can effectively be achieved by swapping memory address offset values or similar techniques."); *see also* Burton, paragraph 79.
[25] Office Action, page 43.

Claim 112

Claim 112 recites a system for executing an artificial neural network. According to pages 45 and 46 of the Office Action, this system is the obvious implementation of Burton's pipelined hardware simulator and Buck's artificial neural network.

Applicant respectfully traverses this rejection because the Office has not considered each and every limitations of claim 112.[26] In this case, page 45 indicates that the CPU 22 in Burton's Figure 6 provides *both the CPU and the processing unit* of claim 112. This is improper; claim 112 recites the CPU and processing unit as separate elements, so they should be considered separately. In addition, the Office Action doesn't identify the controller of claim 112 anywhere in Burton.

Nevertheless, solely in an effort to expedite allowance, Applicant is amending claim 112 to specify that the controller "dictate[s] an order of execution of instructions to the processing unit to perform the computation of the first layer of neurons." Burton doesn't disclose such a controller. Buck doesn't provide such a controller either as explained above with respect to claim 1. Therefore, no combination of Burton and Buck yields the system of claim 112.

7. Other Rejections Under 35 U.S.C. § 103(a)

Claims 8, 9, 17, 18, 44–46, 96, 124, 125, 133, and 134 were rejected as unpatentable over Buck in view of US 7,525,547 to Diard.

Claims 41, 55, and 91 were rejected as unpatentable over Buck in view of US 7,119,810 to Sumanaweera.

Claims 73, 78, 84, 89, and 139 were rejected as unpatentable over Buck in view of US 2006/0129506 to Edelman.

Applicant respectfully traverses these rejections. In any event, as each of the above-identified claims depend from independent claims which are in patentable condition, these rejections are moot.

8. General Comments on Dependent Claims

Since each of the dependent claims depends from a base claim that is believed to be in condition for allowance, Applicant believes that it is unnecessary at this time to argue the

[26] MPEP 2143.03 ("Examiners must consider all claim limitations when determining patentability of an invention over the prior art.").

allowability of each of the dependent claims individually. However, Applicant does not necessarily concur with the interpretation of the dependent claims as set forth in the Office Action, nor does Applicant concur that the basis for the rejection of any of the dependent claims is proper. Therefore, Applicant reserves the right to specifically address the patentability of the dependent claims in the future.

CONCLUSION

It is respectfully believed that all of the rejections set forth in the Office Action have been addressed. However, the absence of a reply to a specific rejection, objection, or comment set forth in the Office Action does not signify agreement with or concession of that rejection, objection, or comment. In addition, there may be reasons for patentability of any or all pending claims (or other claims) that have not been expressed. Furthermore, nothing in this paper should be construed as an intent to concede any issue with regard to any claim.

Applicant believes that the present application is now in condition for allowance. Favorable reconsideration of the application as amended is respectfully requested.

The Examiner is invited to contact the undersigned by telephone if it is felt that a telephone interview would advance the prosecution of the present application.

The Director is hereby authorized to charge any appropriate fees under 37 C.F.R. §§ 1.16, 1.17, and 1.21 that may be required by this paper and to credit any overpayment, to Deposit Account No. 60-1876. If an extension of time is required to make this reply timely, please charge the appropriate extension of time fee under 37 C.F.R. § 1.17 to Deposit Account No. 60-1876 if not already paid via EFS.

Dated: September 13, 2019

Respectfully submitted,

Smith Baluch LLP
376 Boylston St, Suite 401
Boston MA 02116

Phone: (617) 947-7280

By /Christopher Max Colice/
Christopher Max Colice
Registration No. 65,634

Electronic Patent Application Fee Transmittal

Application Number:	15808201
Filing Date:	09-Nov-2017
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Filer:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
CLAIMS IN EXCESS OF 20	2202	7	50	350
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension-of-Time:				
Extension - 1 month with $0 paid	2251	1	100	100
Miscellaneous:				
			Total in USD ($)	**450**

Electronic Acknowledgement Receipt

EFS ID:	37153298
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice/Lindsey Lovato
Filer Authorized By:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03
Receipt Date:	13-SEP-2019
Filing Date:	09-NOV-2017
Time Stamp:	12:56:41
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$450
RAM confirmation Number	E20199CC57064052
Deposit Account	601876
Authorized User	Lindsey Lovato

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

37 CFR 1.16 (National application filing, search, and examination fees)

37 CFR 1.17 (Patent application and reexamination processing fees)

37 CFR 1.19 (Document supply fees)

37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	NRLA-001US03_IDSTransmittal_20190913.pdf	112546 771052086b74c55f799f4464bd8c29082aea132c	no	1
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
2	Non Patent Literature	Brain_NPL.pdf	699850 875ab8c14f1609c6ca4e2a4e956fd7fa8037153f	no	1
Warnings:					
Information:					
3		NRLA-001US03_Amendment_20190913.pdf	504085 75b843b72a608f3c926541ab23a2a475dfde9877	yes	43

Multipart Description/PDF files in .zip description

Document Description	Start	End
Response After Final Action	1	1
Claims	2	25
Applicant Arguments/Remarks Made in an Amendment	26	43

Warnings:					
Information:					
4	Fee Worksheet (SB06)	fee-info.pdf	32256 596b550d54de92231fd8a8d7d9a9d2a4373111f1	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	1348737		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Attorney Docket No.: NRLA-001US03 **PATENT**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re Application of: Anatoli Gorchechnikov	Confirmation No.: 1077
Application No.: 15/808,201	Group Art Unit: 3992
Filed: 2017-11-09	Examiner: William H WOOD

FOR: **GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD**

VIA EFS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

INFORMATION DISCLOSURE STATEMENT UNDER 37 C.F.R. §1.97(f)

The undersigned respectfully submits this information disclosure statement (IDS) under 37 C.F.R. 1.97(f) and requests consideration of the following reference:

Versace, Brain-inspired computing. Invited keynote address, Bionetics 2010, Boston, MA, USA. 1 page.

Applicant made a bona fide attempt to cite this item in an IDS on March 20, 2019, but inadvertently submitted a copy the wrong reference.

If a fee is due under 37 C.F.R. 1.97(f), the Director is hereby authorized to charge that fee to Deposit Account No. 60-1876 from which the undersigned is authorized to draw.

Dated: September 13, 2019

Respectfully submitted,
SMITH BALUCH LLP

SMITH BALUCH LLP
376 Boylston St., Suite 401
Boston, MA 02116

Tel: (617) 947-7280

By: /Christopher Max Colice/
Christopher Max Colice
Reg. No. 65,634

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 15/808,201	Filing Date 11/09/2017	☐ To be Mailed

ENTITY: ☐ LARGE ☑ SMALL ☐ MICRO

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A		N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A		N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*		x $40 =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*		x $210 =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).				
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))					
* If the difference in column 1 is less than zero, enter "0" in column 2.				TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT	09/13/2019	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 67	Minus	** 139	= 0		x $50 =	0
	Independent (37 CFR 1.16(h))	* 6	Minus	*** 16	= 0		x $230 =	0
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=		x $0 =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=		x $0 =	
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.	LIE
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".	/PATRICIA F LEWIS/
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".	
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.	

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 10/09/2019

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

NOTIFICATION DATE	DELIVERY MODE
10/09/2019	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docketing@smithbaluch.com
eofficeaction@appcoll.com
smith@smithbaluch.com

Advisory Action Before the Filing of an Appeal Brief	Application No. 15/808,201	Applicant(s) Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	AIA (FITF) Status No

--The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

THE REPLY FILED 13 September 2019 FAILS TO PLACE THIS APPLICATION IN CONDITION FOR ALLOWANCE.

NO NOTICE OF APPEAL FILED

1. ☑ The reply was filed after a final rejection. No Notice of Appeal has been filed. To avoid abandonment of this application, applicant must timely file one of the following replies: (1) an amendment, affidavit, or other evidence, which places the application in condition for allowance; (2) a Notice of Appeal (with appeal fee) in compliance with 37 CFR 41.31; or (3) a Request for Continued Examination (RCE) in compliance with 37 CFR 1.114 if this is a utility or plant application. Note that RCEs are not permitted in design applications. The reply must be filed within one of the following time periods:
   a) ☑ The period for reply expires 3 months from the mailing date of the final rejection.
   b) ☐ The period for reply expires on: (1) the mailing date of this Advisory Action; or (2) the date set forth in the final rejection, whichever is later. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of the final rejection.
   c) ☐ A prior Advisory Action was mailed more than 3 months after the mailing date of the final rejection in response to a first after-final reply filed within 2 months of the mailing date of the final rejection.The current period for reply expires ____ months from the mailing date of *the prior Advisory Action* or SIX MONTHS from the mailing date of the final rejection, whichever is earlier.
   Examiner Note: If box 1 is checked, check either box (a), (b) or (c). ONLY CHECK BOX (b) WHEN THIS ADVISORY ACTION IS THE FIRST RESPONSE TO APPLICANTS FIRST AFTER-FINAL REPLY WHICH WAS FILED WITHIN TWO MONTHS OF THE FINAL REJECTION. ONLY CHECK BOX (c) IN THE LIMITED SITUATION SET FORTH UNDER BOX (c). See MPEP 706.07(f).

Extensions of time may be obtained under 37 CFR 1.136(a). The date on which the petition under 37 CFR 1.136(a) and the appropriate extension fee have been filed is the date for purposes of determining the period of extension and the corresponding amount of the fee. The appropriate extension fee under 37 CFR 1.17(a) is calculated from: (1) the expiration date of the shortened statutory period for reply originally set in the final Office action; or (2) as set forth in (b) or (c) above, if checked. Any reply received by the Office later than three months after the mailing date of the final rejection, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

NOTICE OF APPEAL

2. ☐ The Notice of Appeal was filed on _____. A brief in compliance with 37 CFR 41.37 must be filed within two months of the date of filing the Notice of Appeal (37 CFR 41.37(a)), or any extension thereof (37 CFR 41.37(e)), to avoid dismissal of the appeal. Since a Notice of Appeal has been filed, any reply must be filed within the time period set forth in 37CFR 41.37(a).

AMENDMENTS

3. ☑ The proposed amendments filed after a final rejection, but prior to the date of filing a brief, will not be entered because
   a) ☑ They raise new issues that would require further consideration and/or search (see NOTE below);
   b) ☐ They raise the issue of new matter (see NOTE below);
   c) ☑ They are not deemed to place the application in better form for appeal by materially reducing or simplifying the issues for appeal; and/or
   d) ☑ They present additional claims without canceling a corresponding number of finally rejected claims.
   NOTE: See Continuation Sheet (See 37CFR 1.116 and 41.33(a)).
4. ☐ The amendments are not in compliance with 37CFR 1.121. See attached Notice of Non-Compliant Amendment (PTOL-324).
5. ☐ Applicants reply has overcome the following rejection(s): _____
6. ☐ Newly proposed or amended claim(s) _____ would be allowable if submitted in a separate, timely filed amendment canceling the non-allowable claim(s).
7. ☑ For purposes of appeal, the proposed amendment(s):(a) ☑ will not be entered, or (b) ☐ will be entered, and an explanation of how the new or amended claims would be rejected is provided below or appended.

AFFIDAVIT OR OTHER EVIDENCE

8. ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____
9. ☐ The affidavit or other evidence filed after final action, but before or on the date of filing a Notice of Appeal will not be entered because applicant failed to provide a showing of good and sufficient reasons why the affidavit or other evidence is necessary and was not earlier presented. See 37 CFR 1.116(e).
10. ☐ The affidavit or other evidence filed after the date of filing the Notice of Appeal, but prior to the date of filing a brief, will not be entered because the affidavit or other evidence failed to overcome all rejections under appeal and/or appellant fails to provide a showing of good and sufficient reasons why it is necessary and was not earlier presented. See 37 CFR 41.33(d)(1).
11. ☐ The affidavit or other evidence is entered. An explanation of the status of the claims after entry is below or attached.

REQUEST FOR RECONSIDERATION/OTHER

12. ☑ The request for reconsideration has been considered but does NOT place the application in condition for allowance because: See Continuation Sheet.
13. ☐ Note the attached Information *Disclosure Statement*(s). (PTO/SB/08) Paper No(s). _____
14. ☐ Other: _____.

STATUS OF CLAIMS

15. The status of the claim(s) is (or will be) as follows:
    Claim(s) allowed:_____.
    Claim(s) objected to:_____.
    Claim(s) rejected:1-60,62-66 and 68-141.
    Claim(s) withdrawn from consideration:_____.

/William H. Wood/ Primary Examiner, Art Unit 3992	conferee: /RSD/

Continuation of 3. NOTE: The numerous amendments to the independent claims would require further consideration and possibly search. The amendments do not place the application in a better form for appeal. Further, the amendments seeks to add additional claim without canceling any claims.

Continuation of REQUEST FOR RECONSIDERATION/OTHER 12. The request for reconsideration has been considered but does NOT place the application in condition for allowance because: The request for reconsideration and presented arguments are directed toward amended claims that would require further consideration and possibly search. As such, the request for reconsideration is not persuasive.

/AJK/

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchechnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT/RESPONSE TO OFFICE ACTION

In response to the Office Action dated May 13, 2019, to which the deadline for responding is September 13, 2019 with a one-month extension of time, Applicant submits the following Amendments and Remarks, and respectfully requests reconsideration of the application in view thereof.

Amendments to the Claims are reflected in the listing of the claims which begins on page 2 of this paper.

Remarks begin on page 26 of this paper.

DO NOT ENTER: /WHW/

PTO/SB/30EFS (02-18)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

REQUEST FOR CONTINUED EXAMINATION(RCE)TRANSMITTAL
(Submitted Only via EFS-Web)

Application Number	15/808,201	Filing Date	2017-11-09	Docket Number (if applicable)	NRLA-001US03	Art Unit	3992
First Named Inventor	Anatoli Gorchechnikov			Examiner Name	William H WOOD		

This is a Request for Continued Examination (RCE) under 37 CFR 1.114 of the above-identified application.
Request for Continued Examination (RCE) practice under 37 CFR 1.114 does not apply to any utility or plant application filed prior to June 8, 1995, or to any design application. The Instruction Sheet for this form is located at WWW.USPTO.GOV

SUBMISSION REQUIRED UNDER 37 CFR 1.114

Note: If the RCE is proper, any previously filed unentered amendments and amendments enclosed with the RCE will be entered in the order in which they were filed unless applicant instructs otherwise. If applicant does not wish to have any previously filed unentered amendment(s) entered, applicant must request non-entry of such amendment(s).

☒ Previously submitted. If a final Office action is outstanding, any amendments filed after the final Office action may be considered as a submission even if this box is not checked.

☐ Consider the arguments in the Appeal Brief or Reply Brief previously filed on _____

☒ Other Amendment after final filed September 13, 2019

☒ Enclosed

☐ Amendment/Reply

☒ Information Disclosure Statement (IDS)

☐ Affidavit(s)/ Declaration(s)

☐ Other _____

MISCELLANEOUS

☐ Suspension of action on the above-identified application is requested under 37 CFR 1.103(c) for a period of months _____
(Period of suspension shall not exceed 3 months; Fee under 37 CFR 1.17(i) required)

☐ Other _____

FEES

☒ **The RCE fee under 37 CFR 1.17(e) is required by 37 CFR 1.114 when the RCE is filed.**
The Director is hereby authorized to charge any underpayment of fees, or credit any overpayments, to Deposit Account No 601876

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT REQUIRED

☒ Patent Practitioner Signature

☐ Applicant Signature

PTO/SB/30EFS (02-18)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Signature of Registered U.S. Patent Practitioner			
Signature	/Nisha Somnath/	Date (YYYY-MM-DD)	2019-10-09
Name	Nisha Somnath	Registration Number	L1194

This collection of information is required by 37 CFR 1.114. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Patent Application Fee Transmittal

Application Number:	15808201
Filing Date:	09-Nov-2017
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Filer:	Nisha Somnath/Holly Johnson
Attorney Docket Number:	NRLA-001US03

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension - 2 months with $100 paid	2252	1	200	200
Miscellaneous:				
RCE- 1ST REQUEST	2801	1	650	650
		Total in USD ($)		**850**

Electronic Acknowledgement Receipt

EFS ID:	37412496
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Nisha Somnath/Holly Johnson
Filer Authorized By:	Nisha Somnath
Attorney Docket Number:	NRLA-001US03
Receipt Date:	09-OCT-2019
Filing Date:	09-NOV-2017
Time Stamp:	17:04:20
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$850
RAM confirmation Number	E201909H04330924
Deposit Account	
Authorized User	

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	NRLA001US03_SB08_10-09-2019.pdf	70862 594542d9a5bfa22159c354dc3439e374f6ea0198	no	4
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
2	Non Patent Literature	Brain-Inspired_Computing_2010_QC.pdf	468160 2110538187c0d39fe353cba4661ee2dad3b39472	no	1
Warnings:					
Information:					
3	Request for Continued Examination (RCE)	NRLA-001US03_USPTO_SB30_RCE_10-09-2019.pdf	468526 8edd54155781b5d0e8e092c930fe1dc21ed6a69d	no	3
Warnings:					
This is not a USPTO supplied RCE SB30 form.					
Information:					
4	Fee Worksheet (SB06)	fee-info.pdf	32502 1292c967268755307e17ec619ada4c641c68da11	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	1040050		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (02-18)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

U.S.PATENTS

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	15808201
	Filing Date	2017-11-09
	First Named Inventor	Anatoli Gorchechnikov
	Art Unit	3992
	Examiner Name	William H WOOD
	Attorney Docket Number	NRLA-001US03

	1	Versace, Brain-inspired computing. Invited keynote address, Bionetics 2010, Boston, MA, USA. 1 page.	☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Nisha Somnath/	Date (YYYY-MM-DD)	2019-10-09
Name/Print	Nisha Somnath	Registration Number	L1194

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 12/16/2019

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

NOTIFICATION DATE	DELIVERY MODE
12/16/2019	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docketing@smithbaluch.com
eofficeaction@appcoll.com
smith@smithbaluch.com

Office Action Summary	Application No. 15/808,201	Applicant(s) Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	AIA (FITF) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☑ Responsive to communication(s) filed on 10/09/2019.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☐ This action is **FINAL.** 2b) ☑ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5) ☑ Claim(s) 1-60,62-66 and 68-148 is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6) ☐ Claim(s) _____ is/are allowed.

7) ☑ Claim(s) 1-60,62-66 and 68-148 is/are rejected.

8) ☐ Claim(s) _____ is/are objected to.

9) ☐ Claim(s) _____ are subject to restriction and/or election requirement

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to **PPHfeedback@uspto.gov.**

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☑ Notice of References Cited (PTO-892)

2) ☑ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date 09/13/2019; 10/09/2019.

3) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date _____.

4) ☐ Other: _____.

The present application is being examined under the pre-AIA first to invent provisions.

DETAILED ACTION

For reissue applications filed before September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the law and rules in effect on September 15, 2012. Where specifically designated, these are "pre-AIA" provisions.

For reissue applications filed on or after September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the current provisions.

Claims 1-60, 62-66, and 68-148 are currently pending in the application. Claims 1-20 are original claims to patent US 9,189,828 B2 and claims 21-60, 62-66, and 68-148 are newly added claims.

Continued Examination Under 37 CFR 1.114

A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 09/13/2019 has been entered.

Reissue

Applicant is reminded of the continuing obligation under 37 CFR 1.178(b), to timely apprise the Office of any prior or concurrent proceeding in which Patent No. 9,189,828 is or was involved. These proceedings would include interferences, reissues, reexaminations, and litigation.

Applicant is further reminded of the continuing obligation under 37 CFR 1.56, to timely apprise the Office of any information which is material to patentability of the claims under consideration in this reissue application.

These obligations rest with each individual associated with the filing and prosecution of this application for reissue. See also MPEP §§ 1404, 1442.01 and 1442.04.

Applicant is notified that any subsequent amendment to the specification and/or claims must comply with 37 CFR 1.173(b).

Information Disclosure Statement

The information disclosure statements (IDS) submitted on 09/13/2019 and 10/09/2019 have been considered in accordance with MPEP 609, including timing and content requirements (see also 37 C.F.R. 1.97 and 1.98). Documents not meeting a particular criteria are lined through and not considered.

Claim Rejections - 35 USC § 112

The following is a quotation of 35 U.S.C. 112(b):
(b) CONCLUSION.—The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the inventor or a joint inventor regards as the invention.

The following is a quotation of 35 U.S.C. 112 (pre-AIA), second paragraph:
The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claims 62-66 and 101-107 are rejected under 35 U.S.C. 112(b) or 35 U.S.C. 112 (pre-AIA), second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which the inventor or a joint inventor, or for pre-AIA the applicant regards as the invention.

Claim 62 recites the limitation "the textures and shaders" in line 7. There is insufficient antecedent basis for this limitation in the claim.

Claim 101 recites the limitation "the second computational cycle" and "the first computational cycle" in the first "storing" limitation. There is insufficient antecedent basis for this limitation in the claim.

Claim 101 is additionally unclear for stating "storing … first output data during the second computational cycle, the first output data generated for a first variable during the first computational cycle … the first output data representing an output of a first neuron" and then later "making the first output data in the first partition of the memory accessible to other neurons in the artificial neural network during the second computational cycle". It is unclear how both "storing" and making available can be applied to the first output data during the second computational cycle. As such, storing the first output data during the second computational cycle is interpreted as already stored first output data that is continuing to be stored during the second computational data. In this case, "storing" is not the initial act of storing, but the act of continuing to currently be stored.

Claim Rejections - 35 USC § 102 and § 103

In the event the determination of the status of the application as subject to AIA 35 U.S.C. 102 and 103 (or as subject to pre-AIA 35 U.S.C. 102 and 103) is incorrect, any correction of the statutory basis for the rejection will not be considered a new ground of rejection if the prior art relied upon, and the rationale supporting the rejection, would be the same under either status.

The following is a quotation of the appropriate paragraphs of pre-AIA 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (a) the invention was known or used by others in this country, or patented or described in a printed publication in this or a foreign country, before the invention thereof by the applicant for a patent.
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

The following is a quotation of pre-AIA 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claim(s) 101-103 and 106-107 is/are rejected under pre-AIA 35 U.S.C. 102(e) as being anticipated by US 7,219,085 B2 *(herein **Buck**)*.

Claim 101

Buck discloses a method for performing a numerical simulation representing an artificial neural network over a plurality of computational cycles *(Buck: column 3, lines 53-65; column 10, lines 43-53; column 11, lines 42-45; column 13, line 25 to column 14, line 15; column 21, lines 60-67; column 23, lines 10-14)*, the method comprising:

storing, in a first partition of a memory *(Buck: column 15, line 46 to column 16, line 59, Buck uses multiple partitions of memory for various sources/destinations of input and output)*, first output data during the second computational cycle *(Buck: column 11, lines 29-41, W updated from previous cycle and currently stored during the second cycle)*, the first output data generated for a first variable during the first computational cycle and being referenced by a first pointer *(Buck: figure 5, elements 505, 525,*

and 230; and column 16, lines 3-59), the first output data representing an output of a first neuron in the artificial neural network *(Buck: figure 7; column 22, lines 1-24)*;

performing, by a processing unit operably coupled to the memory *(Buck: figure 2, element 220)*, at least one calculation on the first output data during the second computational cycle *(Buck: column 11, lines 29-41; column 12, lines 16-26, W continually output over multiple loops/cycles)*, the at least one calculation yielding second output data representing an output of the first neuron in the artificial neural network *(Buck: figure 7; column 22, lines 1-24)*;

making the first output data in the first partition of the memory accessible to other neurons in the artificial neural network during the second computational cycle *(Buck: column 9, lines 47-61; column 22, lines 1-24)*;

storing, in a second partition of the memory, the second output data, the second output data being referenced by a second pointer *(Buck: column 15, line 46 to column 16, line 59, Buck uses multiple partitions of memory for various sources/destinations of input and output)*; and

swapping the first pointer and the second pointer at the conclusion of the second computational cycle such that the second output data becomes an input for a third computational cycle of the plurality of computational cycles *(Buck: column 15, line 46 to column 16, line 60)*.

Claim 102

Buck discloses the method of claim 101, further comprising: transferring input data into the memory from another memory coupled to a central processing unit (CPU) before the first computational cycle *(Buck: column 10, lines 36-43)*.

Claim 103

Buck discloses the method of claim 101, further comprising: providing, to a central processing unit (CPU), first output data generated during the first computational cycle after a conclusion of the first computational cycle *(Buck: column 11, lines 37-41)*.

Claim 106

Buck discloses the method of claim 101, further comprising: dictating, to the processing unit, an order of execution of instructions *(Buck: column 11, lines 42-58; column 14, lines 9-66)*.

Claim 107

Buck discloses the method of claim 101, wherein performing the at least one calculation comprises receiving, by the processing unit, at least the portion of the input data and a series of instructions and executing, by the processing unit, the instructions on the at least the portion of the input data *(Buck: column 10, lines 36-50; figure 5)*.

Claims 1-11, 21-40, 42-54, 56-60, 62-66, 68-72, 74-77, 79-83, 85-88, 115, 117-136, 137-138, 140-146, and 148 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,477,256 B1 *(herein **Johnson**)* in further view of US 7,525,547 B1 *(herein **Diard**)*.

Claim 1

Buck discloses a computer system *(Buck: figure 1)*, comprising:

a central processing unit to receive input data *(Buck: figure 1, elements 120 and 160; figure 2, element 120; column 9, lines 33-36)*;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit *(Buck: figure 1, elements 121, 130, 141, 152, 156)*;

an accelerator, operably coupled to the central processing unit and the main memory via the bus *(Buck: figure 1, element 190)*, to receive at least a portion of the input data from the main memory *(Buck: column 21, lines 54-64)*, the accelerator comprising:

at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, intermediate computations in the sequence of computations yielding intermediate results *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)*; and

accelerator memory, operably coupled to the at least one graphics processing unit, to store the results of the sequence of computations *(Buck: figures 2 and 5, element 230)*; and

"to initialize textures and shaders in the accelerator memory for performing the sequence of computations, to control performance of the sequence of computations by the at least one graphics processing unit" *(Buck: column 13, line 23 to column 21, line 34)*.

Johnson demonstrates, to the extent *Buck* does not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the at least one graphics processing unit and the accelerator memory, to transfer the at least a portion of the input data into the accelerator memory, and to transfer at least a portion of the output data from the accelerator memory to the main memory" *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. *Buck* demonstrates "to initialize textures and shaders in the accelerator memory for performing the sequence of computations, to control performance of the sequence of computations by the at least one graphics processing unit" *(Buck: column 13, line 23 to column 21, line 34)*. It would have been obvious to one of ordinary skill in

the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data and thus initializing and controlling performance of *Buck*'s GPU, shaders, and textures as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly state "to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one graphics processing unit", ***Diard*** demonstrates that it was known at the time of invention to transfer between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring input and output during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with output data transferred from the accelerator memory to the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 2

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the central processing unit is configured to receive the input data in response to a user interaction *(Buck: column 7, lines 51-55; column 23, lines 13-15)*.

Claim 3

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein: the central processing unit is configured to receive the input data at a first rate; and the at least one graphics processing unit is configured to perform the sequence of computations at a second rate different than the first rate *(Buck: column 2, lines 55-67; column 9, lines 10-14; column 14, lines 40-44; CPU and GPU function at different rates)*.

Claim 4

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory *(Buck: column 7, lines 39-51; figure 5, elements 580 and 585)*.

Claim 5

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*.

Claim 6

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein accelerator memory comprises: a first memory bank to store parameters common to all of the computations in the sequence of computations; and a second memory bank to store data specific to at least one computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Claim 7

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W; figure 5, temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W)*.

Claim 8

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one graphics processing unit has begun to perform another sequence of computations *(Buck: column 12, lines 16-26)*.

Claim 9

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 8, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other

sequence of computations by the at least one graphics processing unit *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*.

Claim 10

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the controller is configured to control execution of the sequence of computations by the at least one graphics processing unit *(in view of Johnson, Buck: column 12, lines 16-26)*.

Claim 11

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, further comprising: at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 23, lines 26-30)*.

Claim 21

Buck discloses an accelerator system *(Buck: figure 2, element 190)* comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, the sequence of computations representing an artificial neural network, computations in the sequence of computations representing layers of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*;

a memory, operably coupled to the GPU, to store the input data and the output data *(Buck: figure 5, element 230)*; and

wherein each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory *(Buck: column 15, lines 46-62)*, the first texture holding a first value computed by the neuron during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture holding a second value computed by the neuron during the current time step of the sequence of computations *(Buck: column 16, lines 3-59) (Buck: generally column 13, line 23 to column 21, line 34)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the GPU and the memory, to transfer the input data into the memory and to transfer at least a portion of the output data from the memory" *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly state "to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the GPU", ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring input

and output during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with output data transferred from the accelerator memory to the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 22

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the GPU is operably coupled to a central processing unit (CPU) of a computer system *(Buck: figure 2)*.

Claim 23

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the GPU is mounted on an expansion card operably coupled to a motherboard of a computer system *(Buck: figure 2, element 210; column 9, lines 15-22, AGP and PCI Express busses are for expansion cards coupled to a motherboard; Johnson: figure 3B)*.

Claim 24

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, further comprising: a central processing unit (CPU) *(Buck: figure 1, element 120)*, a system memory *(Buck: figure 1, element 130)*, and

a non-volatile memory installed on a motherboard *(Buck: figure 1, element 133; column 6, lines 27-40, at least ISA bus indicates motherboard; Johnson: figure 3B)*.

Claim 25

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 24, wherein the GPU is operably coupled to the CPU via a system bus on the motherboard *(Buck: figure 1, element 190 containing the GPU; figure 2, elements 190 and 210; column 9, lines 15-22, AGP and PCI Express busses are for expansion cards coupled to a motherboard; Johnson: figure 3B)*.

Claim 26

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 25, wherein the system bus has a bandwidth of 3.2 GB/s to 15 GB/s *(Buck: column 9, lines 15-22, PCI Express bus has a bandwidth higher than 3.2GB/s; Johnson: column 3, lines 51-65)*.

Claim 27

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 25, wherein the controller is configured to transfer the input data and the at least a portion of the output data via the system bus *(Buck: figure 2, element 210)*.

Claim 28

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 25, further comprising: a video card, operably coupled to the motherboard via a local bus, to transfer graphical information to a display *(Buck: figure 1, elements 190, 121, and 191; claim language "local bus" is reference language "system bus")*.

Claim 29

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 28, wherein the local bus has a lower bandwidth than the system bus *(Buck: figure 1, element 121; column 6, lines 27-40, at least ISA bus; claim language "local bus" is reference language "system bus"; figure 2, element 210; column 9, lines 15-22; ISA bus has a lower bandwidth than a PCI Express bus)*.

Claim 30

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the memory comprises: a first memory partition to store parameters common to every computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4; see also figure 9 and column 15, line 46 to column 16, line 59)*.

Claim 31

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 30, wherein the memory further comprises: a second memory partition to store data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4; see also figure 9 and column 15, line 46 to column 16, line 59)*.

Claim 32

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 31, wherein the second memory partition comprises: a first partition to hold external input data patterns *(Buck: figure 5, element 530; column 10, lines 61-62)*; a second partition to hold representations of internal variables *(Buck: column*

15, lines 49-60); a third partition to hold an input of the computation in the sequence of computations *(Buck: column 10, lines 43-44)*; and a fourth partition to hold the output of the computation in the sequence of computations *(Buck: figure 5, element 510)*.

Claim 33

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 31, wherein the second memory partition is configured to accumulate the output of the computation in the sequence of computations during execution, by the GPU, of computational elements of the computation in the sequence of computations *(Buck: column 11, lines 50-52)*.

Claim 34

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 31, wherein the second memory partition is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to computational elements executed by the GPU during the computation in the sequence of computations *(Buck: figure 5, at least element 510; column 11, line 12 to column 12, line 15)*.

Claim 35

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the computation in the sequence of computations performed by the GPU comprises a plurality of computational elements *(Buck: column 21, lines 61-64)*.

Claim 36

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 35, wherein the plurality of computational elements represents a layer of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27; output of a neural network is an output layer of neurons)*.

Claim 37

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 36, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 38

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the memory is configured to store intermediate results of the sequence of computations *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)* and the controller is configured to transfer the output data from the memory to a main memory without transferring the intermediate results *(Buck: figure 5, temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W)*.

Claim 39

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the controller is configured to initiate transfer of at least a portion of the input data *(Buck: column 10, lines 36-43)* and to transfer the at least a portion of the output data in parallel *(figure 5, element 585; and column 6, lines 35-40, at least ISA bus is 8-bit wide indicating at least a portion of data is in parallel)*.

Claim 40

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the input data is acquired by at least one of a video camera, a microphone, or a cell recording electrode *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64)*.

Claim 42

Buck discloses a method of performing a sequence of computations representing an artificial neural network *(Buck: column 21, lines 61-64)*, the method comprising:

receiving, at a central processing unit (CPU), first input data acquired from an external system *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15; external systems ranging from input devices to the GPU)*;

initializing textures and shaders in a memory operably coupled to the GPU *(Buck: column 15, line column 14, line 17 to column 16, lines 60)*;

transferring the first input data received by the CPU to the memory operably coupled to the GPU *(Buck: column 11, line 12 to column 12, line 26)*;

performing, by a graphics processing unit (GPU), a first computation in the sequence of computations on first input data based on the textures and shaders to generate first output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, computations in the sequence of computations representing respective layers of neurons in the artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*, an output of a computation in the sequence of computations representing an output of a neuron in the artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*;

storing, in a memory operably coupled to the GPU, the first input data and the first output data *(Buck: figure 5, element 230)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller operably coupled to a GPU *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly state "transferring second input data into the memory during performance of the sequence of computations by the GPU", ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring data during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and controller of *Johnson* with data transferred between accelerator memory and the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 43

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 42, wherein transferring the second input data comprises transferring the second input data via a bus operably coupled to the central processing unit (CPU) *(Buck: figure 2)*.

Claim 44

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 42, further comprising: transferring the first output data from the memory to another memory during a second computation in the sequence of computations. ***Diard*** demonstrates that it was known at the time of invention to transfer output from a memory to another memory during a computation *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 13, lines 4-25)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of ***Buck*** with output transferred from the accelerator memory to the main memory after the GPU has started other computations as suggested by ***Diard***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both ***Buck*** and ***Diard*** are directed to CPU and GPU processing.

Claim 45

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 44, further comprising: storing intermediate results of the sequence of computations in the memory, and wherein transferring the first output data from the memory to the other memory occurs without transferring the intermediate results of the sequence of computations *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions*

of W; figure 5, temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W).

Claim 46

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 44, wherein transferring the second input data and transferring the first output data occurs in parallel *(Buck: figure 5, element 585; and column 6, lines 35-40, at least ISA bus is 8-bit wide indicating at least a portion of data is in parallel; and column 10, lines 33-35, GPU and CPU compute at the same time)*. Also, ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 13, lines 4-25)*. ***Diard*** further demonstrates simultaneous transfer of input data and output data by an accelerator controller *(Diard: column 6, lines 10-26; column 13, lines 4-25)*.

Claim 47

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 42, further comprising: storing, in a first memory partition of the memory, parameters common to all of the computations in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Claim 48

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 47, further comprising: storing, in a second memory partition of the memory, data specific to the first computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Claim 49

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 48, further comprising: storing, in the second memory partition, external input data patterns, representations of internal variables, an input of the computation in the sequence of computations, and the output of the computation in the sequence of computations *(Buck: figure 5, element 530; column 10, lines 61-62; column 15, lines 49-60; column 10, lines 43-44; figure 5, element 510)*.

Claim 50

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 42, wherein storing the first output data comprises: accumulating, in the memory, outputs of computational elements executed by the GPU in performing the first computation in the sequence of computations *(Buck: column 11, lines 50-52)*.

Claim 51

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 42, further comprising: storing, in the memory, an output of a previous computation in the sequence of computations; and accessing, by the GPU, the output of the previous computation during performance of the computation in the sequence of computations *(Buck: figure 5, at least element 510; column 11, line 12 to column 12, line 15)*.

Claim 52

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 42, wherein performing the first computation comprises executing a plurality of computational elements representing a layer of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27; output of a neural network is an output layer of neurons)*.

Claim 53

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 52, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 54

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 42, further comprising: acquiring the input data with at least one of a video camera, a microphone, or a cell recording electrode *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64)*.

Claim 56

Buck discloses an accelerator system *(Buck: figure 2, element 190)* comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network *(Buck: column 21, lines 61-64)*, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*; and

a memory, operably coupled to the GPU, to store parameters common to all of the computations in the sequence of computations and data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4; see also figure 9 and column 15, line 46 to column 16, line 59)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the GPU and the memory, to transfer at least a portion of the input data into the memory" and "to transfer at least a portion of the output data from the memory" *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly state to transfer at least a portion of the data between memory "in parallel with performance of the sequence of computations" by the GPU, ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring input and output during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with data transferred between the accelerator memory and the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a

substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 57

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 58

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein the memory is configured to hold: external input data patterns *(Buck: figure 5, element 530; column 10, lines 61-62)*; representations of internal variables *(Buck: column 15, lines 49-60; see also figure 9 and column 15, line 46 to column 16, line 59)*; an input to the computation in the sequence of computations *(Buck: column 10, lines 43-44)*; and the output of the computation in the sequence of computations *(Buck: figure 5, element 510)*.

Claim 59

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein the memory is configured to accumulate the output of the computation in the sequence of computations during execution of the plurality of computational elements by the GPU *(Buck: column 11, lines 50-52; see also figure 9 and column 15, line 46 to column 16, line 59)*.

Claim 60

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein the memory is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to the plurality of computational elements during the computation in the sequence of computations *(Buck: figure 5, at least element 530; column 11, line 12 to column 12, line 15)*.

Claim 62

Buck discloses a method comprising:

receiving, at a central processing unit (CPU), input data acquired from an external system *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15; external systems ranging from input devices to the GPU)*;

initializing graphics processing unit (GPU) programs for performing a sequence of computations on the input data received by the CPU *(Buck: column 15, line column 14, line 17 to column 16, lines 60)*;

performing, with the GPU, a sequence of computations on input data to generate output data based on the textures and shaders *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, a computation in the sequence of computations comprising a plurality of computational elements, the plurality of computational elements representing a layer of neurons in an artificial neural network *(Buck: column 21, lines 61-64)*, an output of the computation in the sequence of computations representing an output of a neuron in the artificial neural network *(Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*;

controlling parameters common to all of the computations in the sequence of computations *(Buck: column 10, line 18 to column 12, line 26)*;

storing, in the memory operably coupled to the GPU, parameters common to all of the computations in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*;

storing, in the memory, data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller operably coupled to a GPU *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly state transferring input and output data between memories during performance of the sequence of computations by the GPU, ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring data during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to

implement the CPU and GPU system of *Buck* and controller of *Johnson* with data transferred between accelerator memory and the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 63-66

The limitations of claims 63-66 correspond to the limitations of claims 57-61 and as such are rejected in a corresponding manner.

Claim 68

Buck discloses a method of executing computations representing an artificial neural network *(Buck: column 21, lines 54-64)* on a computer system comprising at least one central processing unit (CPU) *(Buck: figure 2, element 120)*, a processing unit *(Buck: figure 2, element 220)*, a first memory partition *(Buck: figure 2, element 230; column 21, lines 46-49, multiple memory partitions/locations; see also column 15, lines 46-60 and figure 9)*, and a second memory partition *(Buck: figure 2, element 230; column 21, lines 46-49, multiple memory partitions/locations; see also column 15, lines 46-60 and figure 9)*, the method comprising:

executing, by the at least one CPU, a user interaction stream *(Buck: column 9, lines 35-36)*, the user interaction stream controlling transfer of inputs to the artificial neural network to the first memory partition and the second memory partition *(Buck: figures 2 and 5; column 10, lines 41-43; column 23, lines 13-20)*;

executing, by the processing unit, a computational stream *(Buck: column 2, lines 24-28; column 8, lines 39-44, the GPU is performing computations)*, representing the artificial neural network *(Buck: column 10, lines 36-50)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller operably coupled to a GPU so that "the computational stream controlling data exchange between the user interaction stream and the computational stream" *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data, and thus the computational stream of the processing unit controlling data exchange between the user interaction stream and the computation stream as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly state exchanging data between the computational stream and the user interaction stream during execution of the computations, ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations/computations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring data during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to

implement the CPU and GPU system of *Buck* and controller of *Johnson* with data exchanged between the user interaction stream and the computational stream during execution of the computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 69

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 68, wherein the user interaction stream controls data exchange between the user interaction stream and the computational stream outside of execution of the computations representing the artificial neural network *(Buck: figure 2, element 130 separated by element 120)*.

Claim 70

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 68, wherein executing the user interaction stream comprises: controlling setting and editing of computational elements of the computations representing the artificial neural network *(Buck: figure 5, see at least 'P')*.

Claim 71

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 68, wherein executing the user interaction stream comprises: controlling setting and editing of parameters of the computations representing the artificial neural network *(Buck: figure 5, see at least 'W', 'L', 'ERRORS')*.

Claim 72

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 68, wherein executing the user interaction stream comprises: controlling setting and editing of parameters of the inputs to the artificial neural network *(Buck: figure 5, see at least 'X'' and 'T')*.

Claim 74

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 68, wherein executing the user interaction stream comprises: parsing elements to be used in the computations representing the artificial neural network *(Buck: column 2, lines 10-14; column 6, lines 10-16; column 10, lines 36-38; column 14, line 19 to column 15, line 18; programming and languages include parsing)*.

Claim 75

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 74, wherein the processing unit comprises a graphics processing unit (GPU) and executing the user interaction stream further comprises: converting the elements into GPU programs *(Buck: column 2, lines 10-14; column 6, lines 10-16; column 10, lines 36-38; column 14, line 19 to column 15, line 18; programming and languages include converting)*.

Claim 76

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 75, wherein executing the user interaction stream comprises: compiling the GPU programs *(Buck: column 2, lines 10-14; column 6, lines 10-16; column 10, lines 36-38; column 14, line 19 to column 15, line 18; DirectX includes compilation)*.

Claim 77

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 76, wherein executing the user interaction stream comprises: transferring the GPU programs to the second memory partition *(Buck: figure 5; and column 10, lines 36-41)*.

Claims 79-83 and 85-87

The limitations of claims 79-83 and 85-87 correspond to the limitations of claims 68-72 and 74-77 and as such are rejected in a corresponding manner.

Claim 88

Buck, ***Johnson*** and ***Diard*** disclose the system of claim 87, wherein the user interaction stream transfers the GPU programs to at least one of the first memory partition or the second memory partition *(Buck: figure 5; column 10, lines 41-43; column 23, lines 13-20; figure 2, element 230; column 21, lines 46-49)*.

Claim 115

To the extent ***Buck*** and ***Johnson*** do not explicitly state the system of claim 112, wherein the controller is further configured to transfer the first output data to another memory during computation of the second layer of neurons *(Buck: column 11, line 12 to column 12, line 26)*, ***Diard*** demonstrates that it was known at the time of invention to transfer data between memories while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring input and output during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; figure 5, inputs and outputs; column 11, line 59 to column 12, line 6; column 23, lines 7-9; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have

been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with output data transferred from the accelerator memory to the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claims 117-127 and 129-136

The limitations of claims 117-127 and 129-136 correspond to the limitations of claims 1-20, 42-43 and 47-54 and as such are rejected in a corresponding manner.

Claim 128

The limitations of claim 128 correspond to the limitations of claims 1-20, 42-43 and 47-54 and as such are rejected in a corresponding manner. Additionally, ***Buck***, ***Johnson*** and ***Diard*** disclose the method of 128 including: "executing a first stream with the CPU, the first stream acquiring input data for the sequence of computations" *(Buck: column 9, lines 35-36)*; and "executing a second stream with the processing unit, the second stream taking priority over the first stream during the sequence of computations" *(Buck: column 2, lines 24-28; column 8, lines 39-44, the GPU is performing computations; and column 10, lines 31-33, CPU blocks)*.

Claim 137

Buck discloses a system for executing an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*, the system comprising:

a first processing unit to receive input data in real time from a sensor *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15)*;

a memory, operably coupled to the first processing unit, to store a first portion of the input data *(Buck: figure 5, element 230)*;

a second processing unit, operably coupled to the memory, to generate first output data from the first portion of the input data, the first output data representing an output of at least one neuron in a first layer of neurons in the artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*

and allow a user to change the input data on the fly during the execution of the artificial neural network *(Buck: column 9, lines 33-36, "interactive use applications, such as speech recognition"; considering Buck, through multiple cycles, routinely exchanges information with the neural network of the GPU, and speech is interactive, that is changing on the fly)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the second processing unit and the memory, to store a second portion of the input data in the memory" *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly transferring data "while the second processing unit generates the first output data", ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring input and output during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; figure 5, inputs and outputs; column 11, line 59 to column 12, line 6; column 23, lines 7-9; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with data transferred between memories while generating data as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 138

Buck, ***Johnson*** and ***Diard*** disclose the system of claim 137, wherein the sensor is a microphone *(Buck: column 7, lines 51-59)*.

Claim 140

Buck, ***Johnson*** and ***Diard*** disclose the system of claim 137, wherein the second processing unit is further configured to generate second output data from the second portion of the input data, the second output data representing an output of at least one neuron in a second layer of neurons in the

artificial neural network *(Buck: figure 5, inputs and outputs; column 11, line 59 to column 12, line 6; figure 7; column 21, line 65 to column 22, line 27)*.

Claim 141

Buck, ***Johnson*** and ***Diard*** disclose the system of claim 140, further comprising: a random access memory, operably coupled to the first processing unit, to receive the first output data from the memory while the second processing unit generates the second output data *(Buck: figure 1, element 132; column 6, lines 53-59; column 7, lines 7-18; column 11, line 59 to column 12, line 6; and column 10, lines 33-35, GPU and CPU compute at the same time)*.

Claim 142

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the controller is configured to inform the central processing unit that the sequence of computations is finished *(in view of the obviousness of claim 1 - Buck: column 12, lines 16-26 and figure 5; and Johnson: column 7, lines 33-55 – show the graphics system including a controller informing the CPU when a sequence is finished)*.

Claim 143

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the controller is configured to reduce a processing load on the central processing unit *(in view of the obviousness of claim 1 - Buck: column 2, line 57 to column 3, line 8; and Johnson: column 7, lines 33-55 – show the graphics system including a controller reducing the load on the CPU)*.

Claim 144

Buck, ***Johnson*** and ***Diard*** disclose the computer system of claim 1, wherein the controller is configured to reduce interactions between the central processing unit and the accelerator *(in view of the obviousness of claim 1 - Buck: column 2, line 57 to column 3, line 8; and Johnson: column 7, lines 33-55 – Buck shows the desirability of reducing interactions)*.

Claim 145

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 68, wherein the computational stream takes priority over the user interaction stream during execution of the computations representing the artificial neural network *(Buck: column 10, lines 31-33, CPU blocks)*.

Claim 146

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 68, wherein the computational stream determines a timing of a transfer of input data from the CPU to the processing unit and a timing of a transfer of output data from the processing unit to the CPU during execution of the computations representing the artificial neural network *(Buck: column 12, lines 16-26; Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*.

Claim 148

Buck, ***Johnson*** and ***Diard*** disclose the method of claim 128, wherein the sequence of computations represents an artificial neural network, each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory *(Buck: column 15, line 46 to column 16, line 59)*, and wherein (A) further comprises, during a current time step of the sequence of computations:

storing a first value of the output variable computed during a previous time step of the sequence of computations in a first portion of the accelerator memory *(Buck: column 22, lines 1-18)*;

making the first value accessible to every neuron in the neural network during a current time step of the sequence of computations *(Buck: column 22, lines 1-18)*; and

accumulating a second value of the output variable computed during the current time step in a second portion of the accelerator memory *(Buck: column 21, lines 8-26)*.

Claims 12-20, 90, 92-100, 104, and 108-111 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,525,547 B1 *(herein **Diard**)*.

Claim 12

Buck discloses a method of performing a sequence of computations on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory *(Buck: figure 1, elements 120 and 160; figure 2, element 120; column 9, lines 33-36; figure 1, elements 121, 130, 141, 152, 156; figure 1, element 190; figure 2, elements 220 and 230)*, the method comprising:

(A) performing, by the GPU, the sequence of computations on a first portion of input data so as to generate a first portion of output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, intermediate computations in the sequence of computations yielding intermediate results *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)*, wherein performing the sequence of computations on the first portion of the input data comprises (i) assigning an output variable to a first texture and a second texture, the output variable being included in a first computational element of a plurality of computational elements *(Buck: figure 7, output elements;*

column 16, lines 3-34), the plurality of computational elements representing the sequence of computations *(Buck: column 13, line 25 to column 14, line 44)* and (ii) accumulating a first value for the output variable in the first texture during a first time step *(Buck: column 16, lines 8-20)*;

(B) transferring a second portion of the input data from the main memory to the accelerator via the bus *(Buck: column 11, line 12 to column 12, line 26)*; and

(C) transferring a second portion of the output data from the accelerator memory to the main memory via the bus *(Buck: column 11, line 12 to column 12, line 26)*; and

(D) performing, by the GPU, the sequence of computations on the second portion of the input data, wherein performing the sequence of computations on the second portion of the input data *(Buck: column 12, lines 16-26)* comprises (i) accumulating a second value for the output variable in the second texture during a second time step *(Buck: column 16, lines 8-20)* and (ii) making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements during the second time step *(Buck: column 15, lines 57-60; column 16, lines 8-20; and column 16, lines 21-34)*.

To the extent *Buck* does not explicitly state "(B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; (C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus", ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring input and output during those sequences of computations *(Buck: column 11, line 12 to*

column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* with data transferred between the accelerator memory and the main memory in parallel with performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claims 13-20

The limitations of claims 13-20 correspond to the limitations of claims 2-7 and 10-11 and as such are rejected in a corresponding manner.

Claims 90, 92-94, 95, and 97-100

The limitations of claims 90, 92-94, 95, and 97-100 correspond to the limitations of claims 1-20 and as such are rejected in a corresponding manner. However, these claims do not require a "controller".

Claim 96

To the extent ***Buck*** does not explicitly state the system of claim 95, wherein the first memory partition and the second memory partition are logical partitions, ***Diard*** demonstrates that it was known at the time of invention for various memory to be shared or logically partitioned between graphics units and CPUs *(Diard: column 4, lines 5-11)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the memory of ***Buck*** as logical partitions as suggested by ***Diard***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would

have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both ***Buck*** and ***Diard*** are directed to CPU and GPU processing.

Claim 104

Diard demonstrates, to the extent ***Buck*** does not explicitly state "transferring the first output data from the memory to another memory during the second computational cycle", that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing cycles of computation and periodically transferring input and output during cycles *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* with data transferred between the accelerator memory and the main memory during the second computational cycle as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 108

Buck discloses a method of performing a numerical simulation representing an artificial neural network *(Buck: column 3, lines 53-65; column 10, lines 43-53; column 11, lines 42-45; column 13, line 25 to column 14, line 15; column 21, lines 60-67; column 23, lines 10-14)*, the method comprising:

performing, by a processing unit, at least one computation during a first computational cycle of the numerical simulation *(Buck: column 11, lines 29-41; column 12, lines 16-26)*, the first computational cycle representing a first layer of neurons in the artificial neural network *(Buck: figure 7; column 22, lines 1-24)*;

storing, by a memory operably coupled to the processing unit, data used to perform the at least one computation *(Buck: figure 5, element 230)*;

controlling performance of the at least one computation by the processing unit *(Buck: figure 5, element 545)*;

controlling transfer of the output data from the processing unit to the memory during the first computational cycle *(Buck: figure 5, element 220 joined to 230)*;

the second computational cycle representing a second layer of neurons in the artificial neural network *(Buck: figure 7; column 22, lines 1-24)*.

Diard demonstrates, to the extent ***Buck*** does not explicitly state "while performing the at least one computation, acquiring input data with a central processing unit and operably coupled to a main memory" and "controlling transfer of input data to the memory from the main memory for use by the processing unit during a second computational cycle of the numerical simulation", that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing cycles of computation and periodically transferring input and output during cycles *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* with data transferred between the accelerator memory and the main memory while the GPU performs

computations, calculations, and computational cycles as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 109

Buck discloses the method of claim 108, further comprising: dictating an order of execution of instructions to the processing unit *(Buck: column 11, lines 42-58; column 14, lines 9-66)*.

Claim 110

Diard demonstrates, to the extent ***Buck*** does not explicitly state "transferring output data from the memory to the main memory while the processing unit performs the at least one computation", that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing cycles of computation and periodically transferring input and output during cycles *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* with data transferred between the accelerator memory and the main memory while the GPU performs at least one computation as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a

predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 111

Buck discloses the method of claim 108, further comprising: performing successive computational cycles of the numerical simulation by feeding output data generated by the processing unit from a previous computational cycle and input data for a next computational cycle into the processing unit *(Buck: column 11, lines 29-41; column 12, lines 16-26)*.

Claims 41 and 55 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,477,256 B1 *(herein **Johnson**)* in view of US 7,525,547 B1 *(herein **Diard**)* in further view of US 7,119,810 B2 *(herein **Sumanaweera**)*.

Claim 41

To the extent ***Buck*** does not explicitly state the accelerator system of claim 21, wherein the input data is loaded from disk. ***Sumanaweera*** demonstrates that it was known at the time of invention to load GPUs from disk *(Sumanaweera: column 1, lines 10-19)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement ***Buck*** with input data loaded from disk as suggested by ***Sumanaweera***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 55

The limitations of claims 55 correspond to the limitations of claim 41 and as such are rejected in a corresponding manner.

Claim 91 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,525,547 B1 *(herein **Diard**)* in further view of US 7,119,810 B2 *(herein **Sumanaweera**)*.

Claim 91

The limitations of claim 91 correspond to the limitations of claim 41 and as such are rejected in a corresponding manner.

Claims 73, 78, 84, 89, and 139 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,477,256 B1 *(herein **Johnson**)* in view of US 7,525,547 B1 *(herein **Diard**)* in further view of US 2006/0129506 A1 *(herein **Edelman**)*.

Claim 73

To the extent, ***Buck*** does not explicitly state the method of claim 68, wherein executing the user interaction stream comprises: specifying an output to be saved to disk and/or displayed on a screen. ***Edelman*** demonstrates that it was known at the time of invention to save computations of an output stream to a disk *(Edelman: paragraphs 0056, and 0079)* and for the output stream to be related to an artificial or simulated neural system *(Edelman: paragraphs 0056, and 0079)*. Further, ***Buck*** demonstrates CPU controlled outputs and disk memory disk *(Buck: figures 1 and 2; column 3, lines 46-51; column 9, lines 15-37)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the artificial neural network of ***Buck*** with saving output to disk as suggested by

Edelman's teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 78

To the extent ***Buck*** does not state the method of claim 68, further comprising: executing, by the at least one CPU, a data output stream, the data output stream controlling transfer of outputs of the computations representing the artificial neural network to disk. ***Edelman*** demonstrates that it was known at the time of invention to save computations of an output stream to a disk *(Edelman: paragraphs 0056, and 0079)* and for the output stream to be related to an artificial or simulated neural system *(Edelman: paragraphs 0056, and 0079)*. Further, ***Buck*** demonstrates CPU controlled outputs and disk memory disk *(Buck: figures 1 and 2; column 3, lines 46-51; column 9, lines 15-37)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the artificial neural network of ***Buck*** with saving output to disk as suggested by ***Edelman***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claims 84 and 89

The limitations of claims 84 and 89 correspond to the limitations of claims 73 and 78 and as such are rejected in a corresponding manner.

Claim 139

Buck does not explicitly state wherein the sensor is a video camera. ***Edelman*** demonstrates that it was known at the time of invention to make use of video camera input sensors *(Edelman: paragraph 0049)* in an environment related to artificial neural networks *(Edelman: paragraph 0006)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement an input sensor of *Buck* with or as a video camera as suggested by *Edelman*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both references are directed to artificial neural networks.

Claim 105 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US Patent 5,142,665 *(herein **Bigus**)*.

Claim 105

To the extent ***Buck*** and ***Johnson*** do not explicitly state the method of claim 101, further comprising: displaying, with a video system, an indication of the numerical simulation. ***Bigus*** demonstrates that it was known at the time of invention to display simulation results *(Bigus: column 3, lines 32-34; column 6, lines 35-41; column 10, lines 29-33)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the simulations of ***Burton*** with displaying an indication of the simulation as suggested by ***Bigus'*** teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claims 112-114 and 116 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,477,256 B1 *(herein **Johnson**)*.

Claim 112

Buck discloses a system for performing a numerical simulation on input data *(Buck: column 3, lines 53-65; column 10, lines 45-47; column 11, lines 42-45; column 13, line 25 to column 14, line 15; column 21, lines 60-64)*, the system comprising:

a central processing unit (CPU) to provide first input data *(Buck: figure 2, element 120; figure 5; column 9, lines 35-36)*;

a memory, operably coupled to the CPU *(Buck: figure 2, element 230)*, to store the first input data in a first partition, referenced by a first pointer *(Buck: figure 5, elements 505, 525, and 230; and column 16, lines 3-59)*, before computing a first layer of neurons of the artificial neural network *(Buck: column 22, lines 1-18)*;

a processing unit, operably coupled to the memory *(Buck: figure 2, element 220)*, to perform, during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data *(Buck: column 22, lines 1-18, see first result H)*, the first output data representing an output of at least one neuron in the first layer of neurons *(Buck: column 22, lines 18-24)*; and

store the first output data in a second partition of the memory, the second partition referenced by a second pointer *(Buck: column 15, line 46 to column 16, line 60)*, and to swap the first pointer with the second pointer at the end of the first computation of the first layer of neurons, such that the first output data becomes an input for a second layer of neurons of the artificial neural network *(Buck: column 15, line 46 to column 16, line 60; column 22, lines 21-23, "because there are 2 layers in this neural network, the same operations are repeated", that is the*

same operations with different inputs, so that the input of the operations must be pointed to the previous output), and

dictate an order of execution of instructions to the processing unit to perform the computation of the first layer of neurons *(Buck: column 11, lines 42-58; column 14, lines 9-66)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the processing unit and the memory, to store" data *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data between memory as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

Claim 113

Buck and ***Johnson*** disclose the system of claim 112, wherein the processing unit comprises a graphics processing unit *(Buck: figure 5)*.

Claim 114

Buck and ***Johnson*** disclose the system of claim 112, wherein the controller is configured to send instructions for performing the at least one calculation to the processing unit *(Buck: column 10, lines 34 to column 12, line 26; and Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*.

Claim 116

Buck and ***Johnson*** disclose the system of claim 112, wherein the memory further comprises: a third partition to store internal variables; and a fourth partition to store data used as input at a particular layer of neurons of the artificial neural network *(Buck: column 15, line 46 to column 16, line 60)*.

Claim 147 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of *(herein **Diard**)* in further view of US 2006/0129506 A1 *(herein **Edelman**)*.

Claim 147

Buck and ***Diard*** disclose the method of claim 108, wherein acquiring the input data comprises acquiring the input data in real time *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 23, lines 26-30)*.

Buck does not explicitly state wherein the sensor is a video camera. ***Edelman*** demonstrates that it was known at the time of invention to make use of video camera input sensors *(Edelman: paragraph 0049)* in an environment related to artificial neural networks *(Edelman: paragraph 0006)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement an input sensor of *Buck* with or as a video camera as suggested by *Edelman*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both references are directed to artificial neural networks.

Response to Arguments

Patent Owner's arguments filed 09/13/2019 (herein: Remarks) have been fully considered but they are not persuasive. Generally, Patent Owner's arguments with respect to claims 1-60, 62-66, and 68-148 have been considered but are moot because the arguments do not apply to the prior art as constituted under the current rejection. However, Patent Owner arguments with regard to claim 137 that *Buck* does not show real-time or "allow a user to change the input data on the fly during the execution of the artificial neural network" is unpersuasive. *Buck* discloses "in interactive use applications, such as speech recognition" *(Buck: column 9, lines 33-36)*. *Buck* shows cycling data between the CPU and the GPU and speech recognition is interactive, changing on the fly, requiring a "back-and-forth".

Correspondence Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to William H. Wood whose telephone number is (571)272-3736 and fax number is (571)273-3736. The examiner can normally be reached on Monday-Friday 7:30am-4pm. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Alexander Kosowski can be reached on (571)272-3744. The Central Reexamination Unit's (CRU) fax number is (571)273-9900.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR systems, see https://ppair.uspto.gov/epatent/portal/home. For questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/William H. Wood/
Primary Examiner, CRU 3992

Conferees:

/RSD/

/ALEXANDER J KOSOWSKI/
Supervisory Patent Examiner, Art Unit 3992

Notice of References Cited	Application/Control No. 15/808,201	Applicant(s)/Patent Under Reexamination Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-7477256-B1	01-2009	Johnson; Philip B.	G06F3/14	345/501
	B					
	C					
	D					
	E					
	F					
	G					
	H					
	I					
	J					
	K					
	L					
	M					

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

CPC - Searched*		
Symbol	Date	Examiner
G06F 9/5027	10/10/2018	/WHW/
G06T 1/20	10/10/2018	/WHW/
G06T 1/60	10/10/2018	/WHW/
G06N 99/005	10/10/2018	/WHW/
G06N 3/063	10/10/2018	/WHW/
G06F 2209/509	10/10/2018	/WHW/

CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

/William H. Wood/ Primary Examiner, Art Unit 3992	

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

Search Notes		
Search Notes	**Date**	**Examiner**
Prosecution History Review	10/10/2018	/WHW/
IDS Review	10/10/2018	/WHW/
Litigation Search	10/10/2018	/WHW/
EAST BRS Keyword Search	10/10/2018	/WHW/
NPL IEEE Search	10/10/2018	/WHW/
IDS Review	05/03/2019	/WHW/
EAST BRS	05/03/2019	/WHW/
EAST BRS	12/10/2019	/WHW/
IDS Review	12/10/2019	/WHW/

Interference Search			
US Class/CPC Symbol	**US Subclass/CPC Group**	**Date**	**Examiner**

/William H. Wood/ Primary Examiner, Art Unit 3992	

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019						
	1	✓	✓	✓						
	2	✓	✓	✓						
	3	✓	✓	✓						
	4	✓	✓	✓						
	5	✓	✓	✓						
	6	✓	✓	✓						
	7	✓	✓	✓						
	8	✓	✓	✓						
	9	✓	✓	✓						
	10	✓	✓	✓						
	11	✓	✓	✓						
	12	✓	✓	✓						
	13	✓	✓	✓						
	14	✓	✓	✓						
	15	✓	✓	✓						
	16	✓	✓	✓						
	17	✓	✓	✓						
	18	✓	✓	✓						
	19	✓	✓	✓						
	20	✓	✓	✓						
	21	✓	✓	✓						
	22	✓	✓	✓						
	23	✓	✓	✓						
	24	✓	✓	✓						
	25	✓	✓	✓						
	26	✓	✓	✓						
	27	✓	✓	✓						
	28	✓	✓	✓						
	29	✓	✓	✓						
	30	✓	✓	✓						
	31	✓	✓	✓						
	32	✓	✓	✓						
	33	✓	✓	✓						
	34	✓	✓	✓						
	35	✓	✓	✓						
	36	✓	✓	✓						
	37	✓	✓	✓						
	38	✓	✓	✓						
	39	✓	✓	✓						
	40	✓	✓	✓						
	41	✓	✓	✓						
	42	✓	✓	✓						

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019						
	43	✓	✓	✓						
	44	✓	✓	✓						
	45	✓	✓	✓						
	46	✓	✓	✓						
	47	✓	✓	✓						
	48	✓	✓	✓						
	49	✓	✓	✓						
	50	✓	✓	✓						
	51	✓	✓	✓						
	52	✓	✓	✓						
	53	✓	✓	✓						
	54	✓	✓	✓						
	55	✓	✓	✓						
	56	✓	✓	✓						
	57	✓	✓	✓						
	58	✓	✓	✓						
	59	✓	✓	✓						
	60	✓	✓	✓						
	61	✓	-	-						
	62	✓	✓	✓						
	63	✓	✓	✓						
	64	✓	✓	✓						
	65	✓	✓	✓						
	66	✓	✓	✓						
	67	✓	-	-						
	68	✓	✓	✓						
	69	✓	✓	✓						
	70	✓	✓	✓						
	71	✓	✓	✓						
	72	✓	✓	✓						
	73	✓	✓	✓						
	74	✓	✓	✓						
	75	✓	✓	✓						
	76	✓	✓	✓						
	77	✓	✓	✓						
	78	✓	✓	✓						
	79	✓	✓	✓						
	80	✓	✓	✓						
	81	✓	✓	✓						
	82	✓	✓	✓						
	83	✓	✓	✓						
	84	✓	✓	✓						
	85	✓	✓	✓						
	86	✓	✓	✓						
	87	✓	✓	✓						
	88	✓	✓	✓						
	89	✓	✓	✓						
	90	✓	✓	✓						
	91	✓	✓	✓						
	92	✓	✓	✓						
	93	✓	✓	✓						
	94	✓	✓	✓						

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019						
	95	✓	✓	✓						
	96	✓	✓	✓						
	97	✓	✓	✓						
	98	✓	✓	✓						
	99	✓	✓	✓						
	100	✓	✓	✓						
	101	✓	✓	✓						
	102	✓	✓	✓						
	103	✓	✓	✓						
	104	✓	✓	✓						
	105	✓	✓	✓						
	106	✓	✓	✓						
	107	✓	✓	✓						
	108	✓	✓	✓						
	109	✓	✓	✓						
	110	✓	✓	✓						
	111	✓	✓	✓						
	112	✓	✓	✓						
	113	✓	✓	✓						
	114	✓	✓	✓						
	115	✓	✓	✓						
	116	✓	✓	✓						
	117	✓	✓	✓						
	118	✓	✓	✓						
	119	✓	✓	✓						
	120	✓	✓	✓						
	121	✓	✓	✓						
	122	✓	✓	✓						
	123	✓	✓	✓						
	124	✓	✓	✓						
	125	✓	✓	✓						
	126	✓	✓	✓						
	127	✓	✓	✓						
	128	✓	✓	✓						
	129	✓	✓	✓						
	130	✓	✓	✓						
	131	✓	✓	✓						
	132	✓	✓	✓						
	133	✓	✓	✓						
	134	✓	✓	✓						
	135	✓	✓	✓						
	136	✓	✓	✓						
	137		✓	✓						
	138		✓	✓						
	139		✓	✓						
	140		✓	✓						
	141		✓	✓						
	142			✓						
	143			✓						
	144			✓						
	145			✓						
	146			✓						

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019						
	147			✓						
	148			✓						

Attorney Docket No.: NRLA-001US03 **PATENT**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re Application of: Anatoli Gorchechnikov	Confirmation No.: 1077
Application No.: 15/808,201	Group Art Unit: 3992
Filed: 2017-11-09	Examiner: William H WOOD

FOR: **GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD**

VIA EFS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

INFORMATION DISCLOSURE STATEMENT UNDER 37 C.F.R. §1.97(f)

The undersigned respectfully submits this information disclosure statement (IDS) under 37 C.F.R. 1.97(f) and requests consideration of the following reference:

Versace, Brain-inspired computing. Invited keynote address, Bionetics 2010, Boston, MA, USA. 1 page.

Applicant made a bona fide attempt to cite this item in an IDS on March 20, 2019, but inadvertently submitted a copy the wrong reference.

If a fee is due under 37 C.F.R. 1.97(f), the Director is hereby authorized to charge that fee to Deposit Account No. 60-1876 from which the undersigned is authorized to draw.

Dated: September 13, 2019

Respectfully submitted,
SMITH BALUCH LLP

SMITH BALUCH LLP
376 Boylston St., Suite 401
Boston, MA 02116

Tel: (617) 947-7280

By: /Christopher Max Colice/
Christopher Max Colice
Reg. No. 65,634

/William H. Wood/ 12/09/2019

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

PTO/SB/08a (02-18)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

U.S.PATENTS

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button.

U.S.PATENT APPLICATION PUBLICATIONS

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2]	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	15808201
	Filing Date	2017-11-09
	First Named Inventor	Anatoli Gorchechnikov
	Art Unit	3992
	Examiner Name	William H WOOD
	Attorney Docket Number	NRLA-001US03

	1	Versace, Brain-inspired computing. Invited keynote address, Bionetics 2010, Boston, MA, USA. 1 page.	☐

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature	/William H. Wood/	Date Considered	12/09/2019

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	15808201
Filing Date	2017-11-09
First Named Inventor	Anatoli Gorchechnikov
Art Unit	3992
Examiner Name	William H WOOD
Attorney Docket Number	NRLA-001US03

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Nisha Somnath/	Date (YYYY-MM-DD)	2019-10-09
Name/Print	Nisha Somnath	Registration Number	L1194

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /WHW/

OK TO ENTER: /WHW/

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchechnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT/RESPONSE TO OFFICE ACTION

In response to the Office Action dated May 13, 2019, to which the deadline for responding is September 13, 2019 with a one-month extension of time, Applicant submits the following Amendments and Remarks, and respectfully requests reconsideration of the application in view thereof.

Amendments to the Claims are reflected in the listing of the claims which begins on page 2 of this paper.

Remarks begin on page 26 of this paper.

OK TO ENTER: /WHW/

12/10/2019

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S140	1	("7219085").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/11/20 13:55
S141	1	S140 and texture	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/20 13:55
S142	1	("7219085").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/11/21 15:55
S143	1	S142 and controller	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/21 15:55
S144	1	S142 and interface	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/21 15:56
S145	1	S142 and block	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/21 16:36
S146	11531	(graphics adj processing adj unit) and accelerator	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:04
S147	293	(graphics adj processing adj unit) same accelerator same memory same controller	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:05
S148	3943	(graphics adj processing adj unit) and (graphics adj accelerator)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:10
S149	794	S148 and (accelerator with interface)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:16
S150	1061	S148 and (accelerator same interface)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:16
S151	176	S148 and (accelerator near3 interface)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:16
S152	26432	((graphics adj processing adj unit) or GPU) with interface	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:26

S153	6267	((graphics adj processing adj unit) or GPU) near3 interface	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:26
S154	8427	((graphics adj processing adj unit) or GPU) near3 controller	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:27
S155	2602	((graphics adj processing adj unit) or GPU) near3 controller near3 memory	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/11/22 11:28
S156	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/12/02 15:43
S157	1	S156 and sequence	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/12/02 15:43
S158	1	S156 and elements	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/12/02 15:47
S159	1	("7477256").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/12/05 13:58
S160	0	S159 and disk	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/12/05 13:58
S161	0	S159 and drive	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/12/05 13:59
S162	0	S159 and disc	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/12/05 13:59
S163	1	("7219085").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/12/05 15:03
S164	0	S163 and simulation	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/12/05 15:03
S165	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/12/05 15:09
S166	1	S165 and numerical	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2019/12/05 15:09

12/ 10/ 2019 12:23:40 PM
C:\ Users\ wwood\ Documents\ EAST\ Workspaces\ 15 808201.wsp

Doc Code: ECOMM.AUTH/ECOMM.WTDW
Doc Description: Internet Communications Authorized/Internet Communications Authorization Withdrawn

PTO/SB/439 (11-15)

AUTHORIZATION FOR INTERNET COMMUNICATIONS IN A PATENT APPLICATION OR REQUEST TO WITHDRAW AUTHORIZATION FOR INTERNET COMMUNICATIONS		
	Application No.	15/808,201
	Filing Date	2017-11-09
	First Named Inventor	Anatoli Gorchechnikov
	Art Unit	3992
	Examiner Name	William H WOOD
	Practitioner Docket No.	NRLA-001US03

To: Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

I. To authorize permission for Internet Communications.

☑ Recognizing that Internet communications are not secure, I hereby authorize the USPTO to communicate with the undersigned and practitioners in accordance with 37 CFR 1.33 and 37 CFR 1.34 concerning any subject matter of this application via video conferencing, instant messaging, or electronic mail. I understand that a copy of these communications will be made of record in the application file. (MPEP 502.03)

II. To withdraw authorization for Internet Communications.

☐ The authorization given on _____, to the USPTO to communicate with the undersigned and any practitioner in accordance with 37 CFR 1.33 and 37 CFR 1.34 concerning any subject matter of this application via Internet communications such as video conferencing, instant messaging, or electronic mail is hereby withdrawn. I understand that the withdrawal is effective when approved rather than when received.

I am the

☐ applicant.

☑ attorney or agent of record. Registration number 65634.

☐ attorney or agent acting under 37 CFR 1.34. Registration number _____.

/Christopher Max Colice/ — Signature

March 2, 2020 — Date

Christopher Max Colice — Typed or printed name

(617) 947-7280 — Telephone Number

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. Juristic entities must be represented by a patent practitioner (see 37 CFR 1.31, which is applicable to any paper filed on or after September 16, 2012, that is presented on behalf of a juristic entity, regardless of application filing date). Submit multiple forms if more than one signature is required, see below*.

☐ * Total of _____ forms are submitted.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	38735959
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice/Lindsey Lovato
Filer Authorized By:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03
Receipt Date:	02-MAR-2020
Filing Date:	09-NOV-2017
Time Stamp:	11:20:25
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Internet Communications Authorized	NRLA-001US03_USPTO_SB439_20200302_111317.pdf	74933 f597dc90d0b080fe2d9f6576bde7b6110ff8a309	no	2

Warnings:

Information:	
Total Files Size (in bytes):	74933

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchechnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT/RESPONSE TO OFFICE ACTION

In response to the Office Action dated December 16, 2019, to which the deadline for responding is April 16, 2019 with a one-month extension of time, Applicant submits the following Amendments and Remarks, and respectfully requests reconsideration of the application in view thereof.

Amendments to the Claims are reflected in the listing of the claims which begins on page 2 of this paper.

Remarks begin on page 22 of this paper.

IN THE CLAIMS:

Set forth below in ascending order, with status identifiers, is a complete listing of all claims currently under examination. Changes to any amended claims are indicated by [[double brackets]], ~~strikethrough~~, and/or underlining. This listing also reflects any cancellation and/or addition of claims.

1. (Four Times Amended) A computer system, comprising:

a central processing unit to receive input data;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the [first] main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, the sequence of computations representing an artificial neural network, intermediate computations in the sequence of computations representing respective layers of the artificial neural network and yielding intermediate results; and

accelerator memory, operably coupled to the [graphic] at least one graphics processing unit, to store the results of the [plurality of sequential] sequence of computations; and

a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to initialize textures and shaders in the accelerator memory for performing the sequence of computations, to control performance of the sequence of computations by the at least one graphics processing unit, to transfer the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the sequence of computations by the at least one graphics processing unit, and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one [graphic] graphics processing unit.

2. The computer system of claim 1, wherein the central processing unit is configured to receive the input data in response to a user interaction.

3. The computer system of claim 1, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one graphics processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

4. The computer system of claim 1, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

5. The computer system of claim 1, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

6. The computer system of claim 1, wherein accelerator memory comprises:

a first memory bank to store parameters common to all of the computations in the sequence of computations; and

a second memory bank to store data specific to at least one computation in the sequence of computations.

7. The computer system of claim 1, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

8. The computer system of claim 1, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one graphics processing unit has begun to perform another sequence of computations.

9. The computer system of claim 8, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one graphics processing unit.

10. The computer system of claim 1, wherein the controller is configured to control execution of the sequence of computations by the at least one graphics processing unit.

11. (Amended) The computer system of claim 1, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central [processor] processing unit, to acquire the input data in real time.

12. (Thrice Amended) A method of performing a sequence of computations representing an artificial neural network on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory, the method comprising:

(A) performing, by the GPU, the sequence of computations on a first portion of [the] input data so as to generate a first portion of [the] output data, the first portion of the output data representing an output of a neuron in a first layer of the artificial neural network, intermediate computations in the sequence of computations yielding intermediate results, wherein performing the sequence of computations on the first portion of the input data comprises (i) assigning an output variable to a first texture and a second texture, the output variable being included in a first computational element of a plurality of computational elements, the plurality of computational elements representing the sequence of computations and (ii) accumulating a first value for the output variable in the first texture during a first time step;

(B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; [and]

(C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus, the second portion of the output data representing an output of a neuron in a second layer in the artificial neural network; and

(D) performing, by the GPU, the sequence of computations on the second portion of the input data, wherein performing the sequence of computations on the second portion of the input data comprises (i) accumulating a second value for the output variable in the second texture during a second time step and (ii) making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements during the second time step.

13. The method of claim 12, further comprising:

storing the input data in the main memory in response to a user interaction.

14. The method of claim 12, further comprising:

receiving the input data at a first rate; and

wherein (A) comprises performing the sequence of computations at a second rate different than the first rate.

15. Cancel claim 15.

16. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory without transferring any of the intermediate results of the plurality of sequential computations from the accelerator memory to the main memory so as to reduce data transfer via the bus.

17. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory after the GPU has begun to perform another sequence of computations.

18. The method of claim 17, wherein (C) further comprises:

initiating transfer of the second portion of the output data in parallel with performance of at least one computation in the other sequence of computations.

19. The method of claim 12, further comprising:

acquiring the input data in real time with at least one of a video camera, a microphone, or a cell recording electrode operably coupled to the CPU.

20. The method of claim 12, further comprising:

storing parameters common to all of the computations in the sequence of computations in a first memory bank in the accelerator memory; and

storing data specific to at least one computation in the sequence of computations in a second memory bank in the accelerator memory.

21. (New, Twice Amended) An accelerator system comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data, the sequence of computations representing an artificial neural

network, computations in the sequence of computations representing respective layers of neurons in an artificial neural network;

a memory, operably coupled to the GPU, to store the input data and the output data; and

a controller, operably coupled to the GPU and the memory, to transfer the input data into the memory and to transfer at least a portion of the output data from the memory during performance of the sequence of computations by the GPU,

wherein each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture holding a first value computed by the neuron during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture holding a second value computed by the neuron during the current time step of the sequence of computations.

22. (New) The accelerator system of claim 21, wherein the GPU is operably coupled to a central processing unit (CPU) of a computer system.

23. (New) The accelerator system of claim 21, wherein the GPU is mounted on an expansion card operably coupled to a motherboard of a computer system.

24. (New) The accelerator system of claim 21, further comprising:

a central processing unit (CPU), a system memory, and a non-volatile memory installed on a motherboard.

25. (New) The accelerator system of claim 24, wherein the GPU is operably coupled to the CPU via a system bus on the motherboard.

26. (New) The accelerator system of claim 25, wherein the system bus has a bandwidth of 3.2 GB/s to 15 GB/s.

27. (New) The accelerator system of claim 25, wherein the controller is configured to transfer the input data and the at least a portion of the output data via the system bus.

28. (New) The accelerator system of claim 25, further comprising:

a video card, operably coupled to the motherboard via a local bus, to transfer graphical information to a display.

29. (New) The accelerator system of claim 28, wherein the local bus has a lower bandwidth than the system bus.

30. (New) The accelerator system of claim 21, wherein the memory comprises:

a first memory partition to store parameters common to every computation in the sequence of computations.

31. (New) The accelerator system of claim 30, wherein the memory further comprises:

a second memory partition to store data specific to the computation in the sequence of computations.

32. (New) The accelerator system of claim 31, wherein the second memory partition comprises:

a first partition to hold external input data patterns;

a second partition to hold representations of internal variables;

a third partition to hold an input of the computation in the sequence of computations; and

a fourth partition to hold the output of the computation in the sequence of computations.

33. (New) The accelerator system of claim 31, wherein the second memory partition is configured to accumulate the output of the computation in the sequence of computations during execution, by the GPU, of computational elements of the computation in the sequence of computations.

34. (New) The accelerator system of claim 31, wherein the second memory partition is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to computational elements executed by the GPU during the computation in the sequence of computations.

35. (New) The accelerator system of claim 21, wherein the computation in the sequence of computations performed by the GPU comprises a plurality of computational elements.

36. (New) The accelerator system of claim 35, wherein the plurality of computational elements represents a layer of neurons in an artificial neural network.

37. (New) The accelerator system of claim 36, wherein all neurons in the layer of neurons are described by the same equation.

38. (New) The accelerator system of claim 21, wherein the memory is configured to store intermediate results of the sequence of computations and the controller is configured to transfer the output data from the memory to a main memory without transferring the intermediate results.

39. (New) The accelerator system of claim 21, wherein the controller is configured to initiate transfer of at least a portion of the input data and to transfer the at least a portion of the output data in parallel.

40. (New) The accelerator system of claim 21, wherein the input data is acquired by at least one of a video camera, a microphone, or a cell recording electrode.

41. (New) The accelerator system of claim 21, wherein the input data is loaded from disk.

42. (New, Thrice Amended) A method of performing a sequence of computations representing an artificial neural network, the method comprising:

receiving, at a central processing unit (CPU), first input data acquired from an external system in real time;

initializing, by a controller operably coupled to a graphics processing unit (GPU), textures and shaders in a memory operably coupled to the GPU;

transferring the first input data received by the CPU to the memory operably coupled to the GPU;

performing, by the graphics processing unit (GPU), a first computation in the sequence of computations on the first input data based on the textures and shaders to generate first output data, computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of the first computation in the sequence of computations representing an output of a first neuron in a first layer in the artificial neural network;

storing, in the memory operably coupled to the GPU, the first input data and the first output data; and

transferring second input data acquired from the external system in real time into the memory operably coupled to the GPU after the GPU starts the first computation and before the GPU starts a second computation of the sequence of computations, an output of the second computation in the sequence of computations representing an output of a second neuron in a second layer in the artificial neural network.

43. (New, Amended) The method of claim 42, wherein transferring the second input data comprises transferring the second input data via a bus operably coupled to the CPU.

44. (New, Amended) The method of claim 42, further comprising:

transferring the first output data from the memory to another memory during the second computation in the sequence of computations.

45. (New) The method of claim 44, further comprising:

storing intermediate results of the sequence of computations in the memory, and

wherein transferring the first output data from the memory to the other memory occurs without transferring the intermediate results of the sequence of computations.

46. (New) The method of claim 44, wherein transferring the second input data and transferring the first output data occurs in parallel.

47. (New) The method of claim 42, further comprising:

storing, in a first memory partition of the memory, parameters common to all of the computations in the sequence of computations.

48. (New) The method of claim 47, further comprising:

storing, in a second memory partition of the memory, data specific to the first computation in the sequence of computations.

49. (New) The method of claim 48, further comprising:

storing, in the second memory partition, external input data patterns, representations of internal variables, an input of the computation in the sequence of computations, and the output of the computation in the sequence of computations.

50. (New) The method of claim 42, wherein storing the first output data comprises:

accumulating, in the memory, outputs of computational elements executed by the GPU in performing the first computation in the sequence of computations.

51. (New) The method of claim 42, further comprising:

storing, in the memory, an output of a previous computation in the sequence of computations; and

accessing, by the GPU, the output of the previous computation during performance of the computation in the sequence of computations.

52. (New) The method of claim 42, wherein performing the first computation comprises executing a plurality of computational elements representing a layer of neurons in an artificial neural network.

53. (New) The method of claim 52, wherein all neurons in the layer of neurons are described by the same equation.

54. (New, Amended) The method of claim 42, further comprising:

acquiring the second input data with at least one of a video camera, a microphone, or a cell recording electrode.

55. (New, Twice Amended) The method of claim 42, further comprising:

loading the second input data from disk.

56. (New, Thrice Amended) An accelerator system comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data;

a memory, operably coupled to the GPU, to store parameters common to all of the computations in the sequence of computations and data specific to the computation in the sequence of computations; and

a controller, operably coupled to the GPU and the memory, to control performance of the sequence of computations and data exchange by the GPU, to transfer at least a portion of the input data into the memory while the GPU performs the sequence of computations, and to transfer at least a portion of the output data from the memory while the GPU performs the sequence of computations, each computation in the sequence of computations representing a corresponding layer of neurons in an artificial neural network.

57. (New, Amended) The accelerator system of claim 56, wherein all neurons in one of the layers of neurons are described by the same equation.

58. (New) The accelerator system of claim 56, wherein the memory is configured to hold:

external input data patterns;

representations of internal variables;

an input to the computation in the sequence of computations; and

the output of the computation in the sequence of computations.

59. (New, Amended) The accelerator system of claim 56, wherein the memory is configured to accumulate the output of the computation in the sequence of computations during execution of a plurality of computational elements by the GPU.

60. (New, Amended) The accelerator system of claim 56, wherein the memory is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to a plurality of computational elements during the computation in the sequence of computations.

61. Cancel claim 61.

62. (New, Thrice Amended) A method comprising:

receiving, at a central processing unit (CPU), input data acquired from an external system;

initializing, by a controller, graphics processing unit (GPU) programs for performing a sequence of computations on the input data received by the CPU;

performing, with the GPU, the sequence of computations on the input data to generate output data based on textures and shaders, computations in the sequence of computations representing respective layers of neurons in an artificial neural network;

controlling, by the controller, data exchange from the GPU to the CPU and from the CPU to the GPU upon initialization of the GPU programs;

storing, in the memory operably coupled to the GPU, parameters common to all of the computations in the sequence of computations;

storing, in the memory, data specific to the computation in the sequence of computations; and

transferring at least a portion of the input data into the memory and at least a portion of the output data from the memory during performance of the sequence of computations by the GPU.

63. (New, Amended) The method of claim 62, wherein all neurons in one of the layers of neurons are described by the same equation.

64. (New) The method of claim 62, further comprising:

storing, in the memory, external input data patterns;

storing, in the memory, representations of internal variables; and

storing, in the memory, an input of the computation in the sequence of computations.

65. (New) The method of claim 62, further comprising:

accumulating, in the memory, the output of the computation in the sequence of computations during execution of a plurality of computational elements by the GPU.

66. (New) The method of claim 62, further comprising:

storing, in the memory, an output of a previous computation in the sequence of computations; and

accessing, by the GPU, the output of the previous computation during the computation in the sequence of computations.

67. Cancel claim 67.

68. Cancel claim 68.

69. Cancel claim 69.

70. Cancel claim 70.

71. Cancel claim 71.

72. Cancel claim 72.

73. Cancel claim 73.

74. Cancel claim 74.

75. Cancel claim 75.

76. Cancel claim 76.

77. Cancel claim 77.

78. Cancel claim 78.

79. Cancel claim 79.

80. Cancel claim 80.

81. Cancel claim 81.

82. Cancel claim 82.

83. Cancel claim 83.

84. Cancel claim 84.

85. Cancel claim 85.

86. Cancel claim 86.

87. Cancel claim 87.

88. Cancel claim 88.

89. Cancel claim 89.

90. (New, Twice Amended) A method of performing a sequence of computations with a processing unit, the method comprising:

receiving input data;

performing, by the processing unit, the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network, wherein performing the sequence of computations on the first portion of the input data comprises accumulating a first value for an output variable in a first texture by a first computational element;

in parallel with performing the sequence of computations on the first portion of the input data by the processing unit, transferring a second portion of the input data from a first memory partition to a second memory partition;

in parallel with performing the sequence of computations on the first portion of the input data by the processing unit, transferring a second portion of the output data from the second memory partition to the first memory partition; and

performing, by the processing unit, the sequence of computations on the second portion of the input data so as to generate a third portion of output data, wherein performing the sequence of computations on the second portion of the input data comprises accumulating a second value for the output variable in the first texture by the first computational element and accessing the first value of the output variable in the first texture by a second computational element.

91. (New) The method of claim 90, wherein receiving the input data comprises loading the input data from disk.

92. (New) The method of claim 90, wherein receiving the input data comprises receiving the input data from at least one of a video camera, a microphone, or a cell recording electrode.

93. (New) The method of claim 90, wherein receiving the input data occurs at a first rate and performing the sequence of computations occurs at a second rate different than the first rate.

94. (New) The method of claim 90, wherein transferring the second portion of the output data comprises transferring the second portion of the output data from the second memory partition to the first memory partition without transferring any of the intermediate results from the second memory partition to the first memory partition.

95. (New, Thrice Amended) A system for performing a sequence of computations, the system comprising:

a camera to generate input data in real time;

a first memory partition;

a second memory partition operably coupled to the first memory partition; and

a processing unit, operably coupled to the camera, the first memory partition, and the second memory partition, to perform the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network,

wherein the first memory partition is configured to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit,

wherein the second memory partition is configured to transfer a second portion of the output data to the first memory partition in parallel with performance the sequence of computations by the processing unit, and

wherein the sequence of computations represents the artificial neural network, each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture holds a first value of the output variable computed during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture accumulates a second value of the output variable computed during the current time step.

96. (New) The system of claim 95, wherein the first memory partition and the second memory partition are logical partitions.

97. (New) The system of claim 95, wherein the processing unit is comprises a graphics processing unit (GPU).

98. (New) The system of claim 95, wherein the processing unit is configured to receive the input data at a first rate and to perform the sequence of computations at a second rate is different than the first rate.

99. (New) The system of claim 95, wherein the second memory partition is configured to transfer the second portion of the output data to the first memory partition without transferring any of the intermediate results to the first memory partition.

100. Cancel claim 100.

101. Cancel claim 101.

102. Cancel claim 102.

103. Cancel claim 103.

104. Cancel claim 104.

105. Cancel claim 105.

106. Cancel claim 106.

107. Cancel claim 107.

108. (New, Twice Amended) A method of performing a numerical simulation representing an artificial neural network, the method comprising:

performing, by a processing unit, at least one computation during a first computational cycle of the numerical simulation, the first computational cycle representing a first layer of neurons in the artificial neural network;

storing, by a memory operably coupled to the processing unit, data used to perform the at least one computation;

controlling performance of the at least one computation by the processing unit;

controlling transfer of the output data from the processing unit to the memory during the first computational cycle;

while performing the at least one computation, acquiring input data with a central processing unit and operably coupled to a main memory; and

controlling transfer of the input data to the memory from the main memory for use by the processing unit during a second computational cycle of the numerical simulation, the second computational cycle representing a second layer of neurons in the artificial neural network.

109. (New) The method of claim 108, further comprising:

dictating an order of execution of instructions to the processing unit.

110. (New, Amended) The method of claim 108, further comprising:

transferring output data from the memory to the main memory while the processing unit performs the at least one computation.

111. (New, Amended) The method of claim 108, further comprising:

performing successive computational cycles of the numerical simulation by feeding output data generated by the processing unit from a previous computational cycle and input data for a next computational cycle into the processing unit.

112. (New, Thrice Amended) A system for executing an artificial neural network, the system comprising:

a central processing unit (CPU) to provide first input data;

a memory, operably coupled to the CPU, to store the first input data in a first partition, referenced by a first pointer, before computing a first layer of neurons of the artificial neural network;

a processing unit, operably coupled to the memory, to perform, during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data, the first output data representing an output of at least one neuron in the first layer of neurons; and

a controller, operably coupled to the processing unit and the memory, to:

store the first output data in a second partition of the memory, the second partition referenced by a second pointer, and to swap the first pointer with the second pointer at the end of the computation of the first layer of neurons, such that the first output data becomes an input for a second layer of neurons of the artificial neural network,

transfer the first output data to another memory during computation of the second layer of neurons, and

dictate an order of execution of instructions to the processing unit to perform the computation of the first layer of neurons.

113. (New) The system of claim 112, wherein the processing unit comprises a graphics processing unit.

114. (New) The system of claim 112, wherein the controller is configured to send instructions for performing the at least one calculation to the processing unit.

115. Cancel claim 115.

116. (New, Amended) The system of claim 112, wherein the memory further comprises:

a third partition to store internal variables; and

a fourth partition to store data used as input at a particular layer of neurons of the artificial neural network.

117. (New, Thrice Amended) A computer system, comprising:

a central processing unit to receive input data acquired from an external system;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one processing unit to perform a sequence of computations representing an artificial neural network on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations representing layers of the neural network and yielding intermediate results; and

accelerator memory, operably coupled to the at least one processing unit, to store the results of the sequence of computations; and

a controller, operably coupled to the at least one processing unit and the accelerator memory, to control transfer of the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the sequence of computations by the at least one processing unit, to control transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one processing unit, and to control performance of the sequence of computations by the GPU.

118. (New) The computer system of claim 117, wherein the central processing unit is configured to receive the input data in response to a user interaction.

119. (New) The computer system of claim 117, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

120. (New) The computer system of claim 117, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

121. (New) The computer system of claim 117, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

122. (New) The computer system of claim 117, wherein accelerator memory comprises:

a first memory partition to store parameters common to all of the computations in the sequence of computations; and

a second memory partition to store data specific to at least one computation in the sequence of computations.

123. (New) The computer system of claim 117, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

124. (New) The computer system of claim 117, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one processing unit has begun to perform another sequence of computations.

125. (New) The computer system of claim 124, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one processing unit.

126. (New) The computer system of claim 117, wherein the controller is configured to control execution of the sequence of computations by the at least one processing unit.

127. (New, Amended) The computer system of claim 117, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time.

128. Cancel claim 128.

129. Cancel claim 129.

130. Cancel claim 130.

131. Cancel claim 131.

132. Cancel claim 132.

133. Cancel claim 133.

134. Cancel claim 134.

135. Cancel claim 135.

136. Cancel claim 136.

137. (New, Amended) A system for executing an artificial neural network, the system comprising:

a first processing unit to receive input data in real time from a sensor;

a memory, operably coupled to the first processing unit, to store a first portion of the input data;

a second processing unit, operably coupled to the memory, to generate first output data from the first portion of the input data, the first output data representing an output of at least one neuron in a first layer of neurons in the artificial neural network;

a controller, operably coupled to the second processing unit and the memory, to store a second portion of the input data in the memory while the second processing unit generates the first output data and to allow a user to change the input data on the fly during the execution of the artificial neural network.

138. (New) The system of claim 137, wherein the sensor is a microphone.

139. (New) The system of claim 137, wherein the sensor is a video camera.

140. (New) The system of claim 137, wherein the second processing unit is further configured to generate second output data from the second portion of the input data, the second output data representing an output of at least one neuron in a second layer of neurons in the artificial neural network.

141. (New) The system of claim 140, further comprising:

a random access memory, operably coupled to the first processing unit, to receive the first output data from the memory while the second processing unit generates the second output data.

142. (New) The computer system of claim 1, wherein the controller is configured to inform the central processing unit that the sequence of computations is finished.

143. (New) The computer system of claim 1, wherein the controller is configured to reduce a processing load on the central processing unit.

144. (New) The computer system of claim 1, wherein the controller is configured to reduce interactions between the central processing unit and the accelerator.

145. Cancel claim 145.

146. Cancel claim 146.

147. (New) The method of claim 108, wherein acquiring the input data comprises acquiring the input data in real time with a video camera.

148. Cancel claim 148.

REMARKS

In response to the Office Action dated December 16, 2019, in the above-identified application, Applicant respectfully requests reconsideration. All of the issues raised in the Office Action have been carefully considered and are addressed in this response. The application as now presented is believed to be in allowable condition.

1. Status of the Claims and Summary of Amendments

Upon entry of this response, claims 1–14, 16–60, 62–66, 90–99, 108–127, 137–144, and 147 will be pending in this application, of which claims 1, 12, 21, 42, 56, 62, 90, 95, 108, 112, 117, and 137 are independent claims.

Claim 1 is being amended to recite that "the sequence of computations represent[s] an artificial neural network" and that "intermediate computations in the sequence of computations represent[] respective layers of the artificial neural network." Claim 1 is also being amended to specify that the controller is configured "to transfer ... input data into the accelerator memory during performance of the intermediate computations ... by the at least one graphics processing unit, and to transfer ... output data from the accelerator memory to the main memory during performance of the intermediate computations ... by the at least one graphics processing unit." Support for this amendment appears at least in the following sections of the '828 patent: column 5, lines 1–18; column 5, line 61, through column 6, line 17; and in column 8, lines 5–33 and lines 45–60.

Independent claim 12 is being amended to recite that the first and second portions of the output data representing outputs of neurons from first and second layers, respectively, of an artificial neural network. Support for this amendment appears at least in the following sections of the '828 patent: claim 15; column 5, lines 24–26, 63, and 64; and column 9, lines 3–7.

Dependent claim 15 is being canceled.

Independent claim 42 is being amended to recite that the "first input data [is] acquired from an external system *in real time*" and that "an output of the first computation in the sequence of computations represent[s] an output of a first neuron *in a first layer* in the artificial neural network." Claim 42 is also being amended to recite "transferring second input data *acquired from the external system in real time* into the memory ... *after the GPU starts the first computation and before the GPU starts a second computation of the sequence of computations, an output of the second computation in the sequence of computations representing an output of a*

second neuron in a second layer in the artificial neural network." This amendment is for clarity and is supported by the following sections of the '828 patent: column 5, line 61, through column 6, line 17; col. 12, lines 56–62; and FIG. 4 (reference numerals 440, 455, and 475).

Dependent claim 44 is being amended to recite "the second computation" instead of "a second computation" in conformance with independent claim 42 as amended.

Independent claim 56 is being amended for clarity. As amended, claim 56 recites "a controller, operably coupled to the GPU and the memory, to control performance of the sequence of computations and data exchange by the GPU, to transfer at least a portion of the input data into the memory *while the GPU performs the sequence of computations*, and to transfer at least a portion of the output data from the memory *while the GPU performs the sequence of computations, each computation in the sequence of computations representing a corresponding layer of neurons in an artificial neural network*."

Dependent claims 57, 59, and 60 are being amended to conform to independent claim 56 as amended.

Independent claim 62 is being amended for clarity. As amended, claim 62 recites "computations in the sequence of computations representing respective layers of neurons in an artificial neural network" and "controlling, by the controller, data exchange from the GPU to the CPU and from the CPU to the GPU upon initialization of the GPU programs."

Claims 68–89 are being canceled solely to expedite prosecution and without prejudice or disclaimer. Applicant reserves the right to pursue the subject matter of these claims in a continuing application.

Claim 95 is being amended to recite "a camera to generate input data in real time."

Dependent claim 100 is being canceled.

Claims 101–107 are being canceled solely to expedite prosecution and without prejudice or disclaimer. Applicant reserves the right to pursue the subject matter of these claims in a continuing application.

Independent claim 112 is being amended to incorporate the subject matter of dependent claim 115.

Dependent claim 115 is being canceled.

Independent claim 117 is being amended in a fashion similar to independent claim 1.

Independent claim 128 and dependent claims 129–136 are being canceled solely to expedite prosecution and without prejudice or disclaimer. Applicant reserves the right to pursue the subject matter of these claims in a continuing application.

Independent claim 137 is being amended to recite "a controller ... to store a second portion of the input data in the memory while the second processing unit generates the first output data, *to transfer the first output data to the first processing unit before the second processing unit generates second output data representing an output of at least one neuron in a second layer of neurons in the artificial neural network*." Support for this amendment appears at least in column 12, lines 56–62, and FIG. 4 (reference numerals 440, 455, and 475) of the '828 patent.

Dependent claims 145, 146, and 148 are being canceled solely to expedite prosecution and without prejudice or disclaimer. Applicant reserves the right to pursue the subject matter of these claims in a continuing application.

2. Prior and Concurrent Judicial Proceedings

To Applicant's knowledge, there are no prior or concurrent judicial proceedings involving U.S. Patent No. 9,189,828.

3. Rejections Under 35 U.S.C. § 112

Claims 62–66 were rejected as indefinite because there is insufficient antecedent basis for the term "the textures and shaders" in independent claim 62. As indicated above, claim 62 is being amended to recite "textures and shaders," overcoming this rejection.

4. Interview Summary

Applicant thanks Examiner Wood for granting and arranging a video interview on March 4, 2020. During the interview, Applicant's representative Christopher Max Colice explained that when Buck's GPU executes a neural network, it doesn't transfer *outputs* from neurons in the layers in the neural network to the CPU. Instead, during training, Buck's GPU may transfer *trained parameters,* which are settings that are adjusted to improve the neural network's performance, to the CPU. And during testing and in interactive applications, Buck's GPU doesn't transfer anything to the CPU while it is executing the layers of the neural network.

Applicant's representative also pointed out that Buck teaches extensively that transferring data between the GPU and CPU is "expensive" and should be avoided:

- "*Data transfer from the CPU to the GPU is relatively expensive*, so the CPU pre-loads as much as possible onto the GPU before entering a training loop."[1]
- "Again, *transfer from the GPU to CPU is quite expensive*, so there is an advantage to grouping and keeping the size of Y to a minimum."[2]
- "Currently, dataflow from the GPU to the CPU is not optimized since graphic cards are typically designed to send data to the screen but not back to the CPU. As a result, with current architectures, it is preferable to keep all the trainable parameters on the GPU. Preferably, all other data is kept on the CPU and the processing is performed on the GPU."[3]

Buck also teaches that minimizing data transfer between the GPU and CPU is especially important for interactive applications, like speech recognition, which "sometimes require[s] very fast response time,"[4] because too much data transfer can slow down the GPU. Buck explains:

> If data transfer to the GPU takes too much time, the benefit of doing computation on the GPU is decreased. *In other words, if as much data is transferred to the GPU as is needed for the computation, the bottleneck becomes the transfer of data and any net improvement will be nullified.*[5]

In addition, Applicant's representative pointed out neither Johnson nor Diard says anything about transferring output data *from a GPU to a CPU*, possibly because the GPUs in Johnson and Diard are used solely for generating graphics for displays.

Examiner Wood reserved judgment on these arguments and recommended amending the claims to emphasize the distinctions between Applicant's work and Buck.

5. Rejections Under 35 U.S.C. § 103 Over Buck in View of Johnson and Diard

Claims 1–11, 21–40, 42–54, 56–60, 62–66, 115, 117–127, 137, 138, 140–144, and 146 (including independent claims 1, 21, 42, 56, 62, 117, and 137) were rejected as unpatentable over US 7,219,085 B2 to Buck in view of US 7,477,256 B1 to Johnson in further view of US 7,525,547 B1 to Diard. (For brevity, Applicant is addressing only the rejections of the independent claims below.) In each of these rejections, the Office Action alleges that it would have been obvious to combine Buck's computer system with Johnson's controller, which

[1] Buck, 10:41–43 (emphasis added).
[2] Buck, 13:15–17 (emphasis added).
[3] Buck, 11:21–27.
[4] Buck, 8:66.
[5] Buck, 3:3–8 (emphasis added).

allegedly transfers input data into accelerator memory and output data from accelerator memory into main memory, and to make these transfers while executing computations with a GPU as allegedly taught by Diard.

Applicant respectfully submits that these claims are patentable over Buck, Johnson, and Diard because they recite transferring input data from a CPU to a GPU, transferring output data from neurons in a neural network from the GPU to the CPU, or both *while the GPU executes the layers of the neural network*. As explained in greater detail below, Buck discloses transferring neural network output data from the GPU to the CPU but is silent on making those transfers while the GPU performs computations representing the layers of a neural network. Diard and Johnson are silent on transferring output data from the GPU to the CPU and on neural networks, possibly because their GPUs are used to generate and send graphics to a display, not perform neural network computations. Moreover, it would not have been obvious to modify Buck's computer system to make these data transfers, especially for interactive applications like speech recognition, because Buck says that the extra transfers would nullify the benefits of using a GPU to execute the neural network.

Claim 1

Claim 1, which recites a CPU, main memory, GPU, accelerator memory, and controller, was rejected as unpatentable over Buck in view of Johnson and Diard, which both relate to using GPUs for graphics processing. According to pages 7–9 of the Office Action, it would have been obvious to combine Buck's computer system with Johnson's controller, which allegedly transfers input data into accelerator memory and output data from accelerator memory into main memory, and to make these transfers while executing computations with a GPU as allegedly taught by Diard. Applicant respectfully submits that this rejection should be withdrawn for at least four reasons.

First, Johnson's controller does not transfer output data from accelerator memory into main memory as in claim 1. The Office Action cites FIG. 4A (elements 410 and 442) and col. 7, lines 33–55 of Johnson to support the proposition that Johnson discloses a controller that transfers output data from accelerator memory into main memory.[6] These parts of Johnson relate to a graphics processor 400 with an interface controller 410 that is connected to a CPU via a primary connection 442 and a bridge. Johnson says that this interface controller 410 "outputs

[6] Office Action, page 8.

program instructions and data to graphics processing pipeline 420 for processing."[7] The interface controller also "output[s] read and write requests to memory management unit 430 to access memory included in the graphics adapter, such as frame buffer memory."[8] However, this section of Johnson is silent on "transfer[ring] at least a portion of the output data [generated by the GPU] from the accelerator memory to the main memory" as in claim 1.

The rest of Johnson is also silent on transferring output data generated by the GPU from the accelerator memory to the main memory. Instead, Johnson discloses transferring output data from the graphics processor 400 to either a display device or another graphics processor, depending on whether the graphics processor 400 is configured as a master graphics processor or a slave graphics processor.[9] In any event, the graphics processor 400 (and hence the interface controller 410) does not transfer output data to the main memory. This could be because "graphic cards are typically designed to send data to the screen but not back to the CPU."[10] Moreover, there's no reason for Johnson's graphics processor 400 to send output data from to main memory because the output data is just pixel data for showing on a display.

Second, contrary to the assertion on page 9 of the Office Action, Diard does not demonstrate that is was known to transfer output data *from* accelerator memory *to* main memory during performance of the sequence of computations by a GPU as in claim 1. According to page 9 of the Office Action, the following sections of Diard "demonstrate[] that it was known at the time of invention to transfer between an accelerator memory and the main memory while the GPU is performing other operations": Figure 1 (transfer arrows); column 5, lines 1–3; column 6, lines 10–26; and column 13, lines 4–25. Although these sections of Diard disclose data transfer from the CPU to the GPU, they don't disclose transferring output data from the GPU to the CPU. Instead, Diard discloses that the output data from the GPU goes to the display: "The graphics subsystem 130 periodically outputs pixel data for an image from display memory 218 and displayed on display device 150."[11] Again, this could be because "graphic cards are typically designed to send data to the screen but not back to the CPU."[12] And again, there's no reason for Diard's GPU to send output data to the CPU because the output data is simply pixel data for showing on the display device 150.

---

[7] Johnson, 7: 44–45.
[8] Johnson, 7:48–51.
[9] *See* Johnson, 8:36–67.
[10] Buck, 11:22–23.
[11] Diard, 4:10 –12.
[12] Buck, 11:22–23.

Third, none of the cited references disclose transferring input and output data between main memory and accelerator memory *while a GPU performs intermediate computations representing layers of an artificial neural network* as in amended claim 1.

Diard and Johnson are both silent on neural networks, so they don't provide this element of amended claim 1.

Buck discloses training and testing a neural network on a GPU but is silent on transferring data to the GPU *while the GPU performs intermediate computations representing layers of the neural network*. During training, Buck's CPU transfers a batch of pre-processed input data, called training data (X'), to the GPU before the GPU begins performing any computations.[13] The GPU then processes the training data by executing the neural network.[14] Once the GPU is done executing the neural network, the CPU may load new training data, and the GPU may execute the neural network again.[15] At no point, however, does the CPU load training data while the GPU executes the layers of the neural network.

Similarly, Buck is silent on transferring input and output data while the GPU performs intermediate computations during testing or deployment. Buck states simply that the CPU loads, preprocesses, and sends input data to the GPU.[16] Once the preprocessed input data has been sent to the GPU, the GPU executes the neural network.[17] The GPU accumulates output data (results Y) and transfers them to the CPU in batches.[18] Buck is silent on whether this output data transfer occurs while the GPU is performing intermediate computations representing the layers of the neural network. Likewise, Buck is silent on transferring data from the CPU to the GPU while the GPU performs those intermediate computations.

Fourth, it would not have been obvious to modify Buck's computer system to transfer input data to the GPU during training or testing. Buck states repeatedly and explicitly that transferring data between the CPU and GPU is "expensive."[19] Buck states that this is a drawback

---

[13] Buck, 11:29–31 ("Once the training data (X') ***has been loaded on the GPU***, the CPU instructs the GPU to run the various shaders need to process the machine learning technique (box 540 [in FIG. 5 of Buck])."")(emphasis added).
[14] *See* Buck, 11:29–58.
[15] *See* Buck, 12:20–25.
[16] Buck, 12:60–62 ("The patterns X then are preprocessed to X' before sending them from the CPU to the GPU (box 620). The X' then is sent to the GPU (arrow 625).").
[17] *See* Buck, 13:4–22.
[18] *See* Buck, 13:10–17.
[19] Buck, 10:41–43 ("Data transfer from the CPU to the GPU is relatively expensive, so the CPU pre-loads as much as possible onto the GPU before entering a training loop."); *see also id.*, 11:21–27 ("Currently, dataflow from the GPU to the CPU is not optimized since graphic cards are typically designed to send data to the screen but not back to the CPU. As a result, with current architectures, it is preferable to keep all the trainable parameters on the GPU. Preferably, all other data is kept on the CPU and the processing is performed on the GPU."), and 13:15–17 ("Again,

for interactive applications like speech or character recognition,[20] which "sometimes require very fast response time."[21] In fact, to avoid these expensive data transfers, Buck teaches minimizing data transfers between the CPU and GPU. Buck explains:

> If data transfer to the GPU takes too much time, the benefit of doing computation on the GPU is decreased. In other words, if as much data is transferred to the GPU as is needed for the computation, the bottleneck becomes the transfer of data and any net improvement will be nullified.[22]

In light of Buck's teachings that data transfers are expensive and time consuming, a person of ordinary skill in the art would have understood that modifying Buck's computer system to transfer data while the GPU performs intermediate computations would slow the system, eliminating the advantages of using a GPU. Thus, a person of ordinary skill in the art would not have been motivated to modify Buck's computer system to transfer data while the GPU performs intermediate computations representing layers of a neural network.

Thus, the proposed combination of Buck, Johnson, and Diard does not teach or suggest a controller that transfers both input data from a CPU to a GPU and output data from the GPU to the CPU while the GPU performs calculations representing intermediate layers of an artificial network as in amended claim 1. If anything, Buck teaches that such a controller would not have been obvious because the data transfers would nullify the advantages of using a GPU with a CPU.

Claim 21

Like claim 1, claim 21 recites that "computations in the sequence of computations represent[] respective layers of neurons in an artificial neural network" and that the controller "transfer[s] the input data into the memory and to transfer at least a portion of the output data from the memory *during performance of the sequence of computations by the GPU*." Thus, claim 21 is patentably distinct from Buck, Diard, and Johnson for at least the same reasons as claim 1.

---

transfer from the GPU to CPU is quite expensive, so there is an advantage to grouping and keeping the size of Y to a minimum.").

[20] Buck, 3:3–8.

[21] Buck, 8:66.

[22] Buck, 3:3–8.

Claim 42

Claim 42 was rejected as obvious over Buck in view of Diard and Johnson. As indicated above, claim 42 is being amended to recite that the "first input data [is] acquired from an external system *in real time*" and that "an output of the first computation in the sequence of computations represent[s] an output of a first neuron *in a first layer* in the artificial neural network." Claim 42 is also being amended to recite "transferring second input data *acquired from the external system in real time* into the memory ... *after the GPU starts the first computation and before the GPU starts a second computation of the sequence of computations, an output of the second computation in the sequence of computations representing an output of a second neuron in a second layer in the artificial neural network*."

As explained above, the proposed combination of Buck, Diard, and Johnson does not disclose transferring input data to GPU memory while the GPU executes computations representing individual layers of a neural network. Moreover, it would not have been obvious to modify Buck's computer system to transfer input data to GPU memory as in amended claim 42 because Buck teaches that data transfers this frequent would negate the benefits of using the GPU.[23]

Claim 56

Claim 56 was rejected as obvious over Buck in view of Diard and Johnson. As indicated above, claim 56 is being amended for clarity and now recites "a controller, operably coupled to the GPU and the memory, to control performance of the sequence of computations and data exchange by the GPU, to transfer at least a portion of the input data into the memory *while the GPU performs the sequence of computations*, and to transfer at least a portion of the output data from the memory *while the GPU performs the sequence of computations, each computation in the sequence of computations occurring during a corresponding time step and representing a corresponding layer of neurons in an artificial neural network*."

As explained above, the proposed combination of Buck, Diard, and Johnson does not disclose transferring input data to and output data from GPU memory while the GPU executes computations representing individual layers of a neural network. Moreover, it would not have been obvious to modify Buck's computer system to transfer input data to GPU memory as in

[23] Buck, 3:3–8.

claim 56 because Buck teaches that these more frequent data transfers would negate the benefits of using the GPU.[24]

Claim 62

Claim 62 was rejected as obvious over Buck in view of Diard and Johnson. As indicated above, claim 62 is being amended for clarity and now recites "controlling, by [a] controller, data exchange from the GPU to the CPU and from the CPU to the GPU upon initialization of the GPU programs" and "transferring at least a portion of the input data into the memory and at least a portion of the output data from the memory" while the GPU performs a sequence of computations, with "computations in the sequence of computations representing respective layers of neurons in an artificial neural network."

As explained above, the proposed combination of Buck, Diard, and Johnson does not disclose transferring input data to and output data from GPU memory while the GPU executes computations representing layers of a neural network. The proposed combination of Buck, Diard, and Johnson also fail to disclose "controlling, by [a] controller, data exchange from the GPU to the CPU and from the CPU to the GPU upon initialization of the GPU programs," in part because Diard and Johnson are silent on transferring data *from* the GPU to the CPU as explained above with respect to claim 1.

Claim 117

Like claim 1, claim 117 is being amended to recite that "the sequence of computations represent[s] an artificial neural network," that "intermediate computations in the sequence of computations represent[] respective layers of the artificial neural network," and that the controller is configured "to transfer ... input data into the accelerator memory during performance of the intermediate computations ... by the at least one graphics processing unit, and to transfer ... output data from the accelerator memory to the main memory during performance of the intermediate computations ... by the at least one graphics processing unit." Thus, amended claim 117 is patentable over Buck in view of Diard and Johnson for at least the reasons as amended claim 1.

[21] Buck, 3:3–8.

Claim 137

Claim 137 was rejected as obvious over Buck in view of Diard and Johnson. Among other things, page 5 of the Office Action equates Buck's disclosure of speech recognition with a *trained* neural network as providing "a controller ... to allow a user to change the input data on the fly during the execution of [an] artificial neural network."

Claim 137 is being amended to specify that the controller "transfer[s] the first output data [representing an output of at least one neuron in a *first layer* of neurons in the artificial neural network] to the first processing unit before the second processing unit generates second output data representing an output of at least one neuron in a *second layer* of neurons in the artificial neural network." As explained above, the proposed combination of Buck, Diard, and Johnson doesn't yield a computer system with controller that transfers output data from the GPU to the CPU while computing neural network layers. Buck discloses transferring neural network output data from the GPU to the CPU but is silent on making those transfers after the GPU computes a first layer and before the GPU computes a second layer in the neural network. Diard and Johnson are silent on transferring output data from the GPU to the CPU and on neural networks. This makes claim 137 patentable over the proposed combination of Buck, Diard, and Johnson.

6. Rejections Under 35 U.S.C. § 103(a) Over Buck in view of Diard

Claims 12–20, 90, 92–100, and 108–111 (including independent claims 12, 90, 95, 108, and 112) were rejected as unpatentable over Buck in view of Diard. In each rejection, the Office Action acknowledges that Buck does not explicitly disclose " (C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus." To fill this gap, the Office Action cites Diard as allegedly "demonstrate[ing] that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations."

Applicant respectfully traverses these rejections because neither Buck nor Diard discloses or suggests transferring output data from accelerator (GPU) memory to main memory while the GPU performs other computations. Per pages 40 and 41 of the Office Action, Buck does not explicitly disclose transferring output data from accelerator (GPU) memory to main memory *while the GPU performs other computations*. Diard is silent on transferring output data from the GPU to the CPU. Therefore, the proposed combination of Buck and Diard does not yield the claimed subject matter. What is more, it would not have been obvious to modify Buck's

computer system to transfer GPU output data to CPU memory while the GPU performs computations, especially for interactive applications like speech recognition, because Buck says that the extra transfers would nullify the benefits of using a GPU to execute the neural network.

Claim 12

As indicated above, claim 12 is being amended to recite, in relevant part, that "(C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus, the second portion of the output data representing an output of a neuron in a second layer in the artificial neural network." As explained above, Buck is silent on transferring the outputs of different neural network layers from accelerator memory to main memory. Diard is silent on transferring any GPU output from GPU memory to main memory and so fails to cure Buck's deficiency, meaning that the proposed combination of Buck and Diard fails to yield claim 12 as amended.

In addition, Buck teaches that data transfers between the CPU and GPU should be avoided because they generally decrease the benefit of performing computations with the CPU.[25] Thus, Buck teaches explicitly that it would not have been obvious to perform additional data transfers to or from the GPU in Buck's computer system.

Claim 90

Page 41 of the Office Action says that "[t]he limitations of claims 90, 92–94, 95, and 97–100 correspond to the limitations of claims 1–20 and as such are rejected in a corresponding manner." This is not the case. Claim 90, for example, recites at least two steps that don't appear in claims 1–20. These steps are:

> in parallel with performing the sequence of computations on the first portion of the input data by the processing unit, transferring a second portion of the output data from the second memory partition to the first memory partition; and
>
> performing, by the processing unit, the sequence of computations on the second portion of the input data so as to generate a third portion of output data,
>
> wherein performing the sequence of computations on the second portion of the input data comprises accumulating a second value for the output variable in the

[25] Buck, 3:3–8 (emphasis added).

first texture by the first computational element and accessing the first value of the output variable in the first texture by a second computational element.

These limitations were added in Applicant's last response and have not been identified anywhere in the prior art. This failure makes the rejection of claim 90 improper. Accordingly, Applicant respectfully requests withdrawal of the rejection of claim 90 and all claims that depend on claim 90, including claims 92–94.

Claim 95

Page 41 of the Office Action says that "[t]he limitations of claims 90, 92–94, 95, and 97–100 correspond to the limitations of claims 1–20 and as such are rejected in a corresponding manner." Again, this is not the case. Independent claim 95 recites first and second memory partitions that don't appear in claims 1–20.

In any event, Applicant is amending claim 95 to recite "a camera to generate input data in real time." The proposed combination of Buck and Diard lacks a camera that generates input data in real time. Therefore, amended claim 95 is patentable over the proposed combination of Buck and Diard.

Claim 108

Claim 108 recites "performing, with a processing unit, at least one computation during a first computational cycle ... representing a first layer of neurons in [an] artificial neural network." It also recites "while performing the at least one computation, acquiring input data with a central processing unit and operably coupled to a main memory; and controlling transfer of the input data to the [processing unit] memory from the main memory for use by the processing unit during a second computational cycle ... representing a second layer of neurons in the artificial neural network."

As explained above, the proposed combination of Buck and Diard does not acquire input data while calculating a first layer of a neural network with a processing unit. The proposed combination of Buck and Diard also fails to transfer that input data from main memory to the processing unit memory for the processing unit to use when calculating a second layer of the neural network. Therefore, claim 108 is patentably distinct from the proposed combination of Buck and Diard.

Claim 112

Claim 112 is being amended to recite that the controller is configured to "transfer the first output data [representing an output of at least one neuron in the first layer of neurons] to another memory during computation of the second layer of neurons." The proposed combination of Buck and Diard doesn't include a controller that transfers output data representing an output of one or more neurons in a first layer of neurons to another memory during computation of a second layer of neurons. Therefore, claim 112 is patentably distinct from the proposed combination of Buck and Diard.

7. Other Rejections Under 35 U.S.C. § 103(a)

Claims 41 and 55 were rejected as unpatentable over Buck in view of Johnson, Diard, and US 7,119,810 B2 to Sumanaweera.

Claim 91 was rejected as unpatentable over Buck in view of Diard and Sumanaweera.

Claims 73, 78, 84, 89, and 139 were rejected as unpatentable over Buck in view of Johnson, Diard, and US 2006/0129506 A1 to Edelman.

Claims 112–114 and 116 were rejected as unpatentable over Buck in view of Johnson.

Claim 147 was rejected as unpatentable over Buck in view of Diard and Edelman.

Applicant respectfully traverses these rejections. In any event, as each of the above-identified claims depend from independent claims which are in patentable condition, these rejections are moot.

8. General Comments on Dependent Claims

Since each of the dependent claims depends from a base claim that is believed to be in condition for allowance, Applicant believes that it is unnecessary at this time to argue the allowability of each of the dependent claims individually. However, Applicant does not necessarily concur with the interpretation of the dependent claims as set forth in the Office Action, nor does Applicant concur that the basis for the rejection of any of the dependent claims is proper. Therefore, Applicant reserves the right to specifically address the patentability of the dependent claims in the future.

CONCLUSION

It is respectfully believed that all of the rejections set forth in the Office Action have been addressed. However, the absence of a reply to a specific rejection, objection, or comment set

forth in the Office Action does not signify agreement with or concession of that rejection, objection, or comment. In addition, there may be reasons for patentability of any or all pending claims (or other claims) that have not been expressed. Furthermore, nothing in this paper should be construed as an intent to concede any issue with regard to any claim.

Applicant believes that the present application is now in condition for allowance. Favorable reconsideration of the application as amended is respectfully requested.

The Examiner is invited to contact the undersigned by telephone if it is felt that a telephone interview would advance the prosecution of the present application.

The Director is hereby authorized to charge any appropriate fees under 37 C.F.R. §§ 1.16, 1.17, and 1.21 that may be required by this paper and to credit any overpayment, to Deposit Account No. 60-1876. If an extension of time is required to make this reply timely, please charge the appropriate extension of time fee under 37 C.F.R. § 1.17 to Deposit Account No. 60-1876 if not already paid via EFS.

Dated: April 16, 2020

Respectfully submitted,

Smith Baluch LLP
376 Boylston St, Suite 401
Boston MA 02116

Phone: (617) 947-7280

By /Christopher Max Colice/
Christopher Max Colice
Registration No. 65,634

Electronic Patent Application Fee Transmittal

Application Number:	15808201
Filing Date:	09-Nov-2017
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Filer:	Christopher M. Colice/Lindsey Lovato
Attorney Docket Number:	NRLA-001US03

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension - 1 month with $0 paid	2251	1	100	100
Miscellaneous:				
			Total in USD ($)	**100**

Electronic Acknowledgement Receipt

EFS ID:	39172593
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice/Lindsey Lovato
Filer Authorized By:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03
Receipt Date:	16-APR-2020
Filing Date:	09-NOV-2017
Time Stamp:	12:57:39
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$100
RAM confirmation Number	E20204FC58040611
Deposit Account	601876
Authorized User	Lindsey Lovato

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

37 CFR 1.16 (National application filing, search, and examination fees)

37 CFR 1.17 (Patent application and reexamination processing fees)

37 CFR 1.19 (Document supply fees)

37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		NRLA-001US03_Amendment_2020416.pdf	475580 006c84f93a2d051fc534459c72d35b81207987e4	yes	36

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Amendment/Req. Reconsideration-After Non-Final Reject	1	1
	Claims	2	21
	Applicant Arguments/Remarks Made in an Amendment	22	36

Warnings:					
Information:					
2	Fee Worksheet (SB06)	fee-info.pdf	30602 3a2faf0bde564ddedbe19bfff19218c23094daf9	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	506182	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 15/808,201	Filing Date 11/09/2017	☐ To be Mailed

ENTITY: ☐ LARGE ☑ SMALL ☐ MICRO

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A		N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A		N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*		x $40 =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*		x $210 =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).				
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))					
* If the difference in column 1 is less than zero, enter "0" in column 2.				TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT	04/16/2020	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 67	Minus	** 139	= 0		x $50 =	0
	Independent (37 CFR 1.16(h))	* 11	Minus	*** 16	= 0		x $230 =	0
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=		x $0 =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=		x $0 =	
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.	LIE
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".	/PATRICIA F LEWIS/
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".	
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.	

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Application Number * 15/808,201 *	Application/Control No. 15/808,201	Applicant(s)/Patent under Reexamination Gorchetchnikov et al.	
	Examiner WOOD, WILLIAM H	Art Unit 3992	

Document Code - DISQ	Internal Document - DO NOT MAIL

TERMINAL DISCLAIMER	☐ APPROVED	☑ DISAPPROVED
Date Filed: 09 November 2017	This patent is subject to a Terminal Disclaimer	

Approved/Disapproved by:

/JEAN PROCTOR/

Technology Center: OPLC

Telephone: (571)272-1040

The application number of the application being examined (or the number of the patent under ree~am or reissue) is not
identified or incorrect.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 06/05/2020

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

NOTIFICATION DATE	DELIVERY MODE
06/05/2020	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docketing@smithbaluch.com
eofficeaction@appcoll.com
smith@smithbaluch.com

Office Action Summary	Application No. 15/808,201	Applicant(s) Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	AIA (FITF) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☑ Responsive to communication(s) filed on 04/21/2020.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☑ This action is **FINAL.** 2b) ☐ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5) ☑ Claim(s) 1-14,16-60,62-66,90-99,108-114,116-127,137-144 and 147 is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6) ☑ Claim(s) 1-20,42-55,95-99,112-114,116 and 142-144 is/are allowed.

7) ☑ Claim(s) 21-41,56-60,62-66,90-94,108-111,117-127,137-141 and 147 is/are rejected.

8) ☐ Claim(s) _____ is/are objected to.

9) ☐ Claim(s) _____ are subject to restriction and/or election requirement

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to **PPHfeedback@uspto.gov.**

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)

2) ☐ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date _____.

3) ☑ Interview Summary (PTO-413) Paper No(s)/Mail Date 20200504.

4) ☐ Other: _____.

The present application is being examined under the pre-AIA first to invent provisions.

DETAILED ACTION

For reissue applications filed before September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the law and rules in effect on September 15, 2012. Where specifically designated, these are "pre-AIA" provisions.

For reissue applications filed on or after September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the current provisions.

Claims 1-14, 16-60, 62-66, 90-99, 108-114, 116-127, 137-144, and 147 are currently pending in the application. Claims 1-14 and 16-20 are original claims to patent US 9,189,828 B2 and claims 21-60, 62-66, 90-99, 108-114, 116-127, 137-144, and 147 are newly added claims.

Reissue

Applicant is reminded of the continuing obligation under 37 CFR 1.178(b), to timely apprise the Office of any prior or concurrent proceeding in which Patent No. 9,189,828 is or was involved. These proceedings would include interferences, reissues, reexaminations, and litigation.

Applicant is further reminded of the continuing obligation under 37 CFR 1.56, to timely apprise the Office of any information which is material to patentability of the claims under consideration in this reissue application.

These obligations rest with each individual associated with the filing and prosecution of this application for reissue. See also MPEP §§ 1404, 1442.01 and 1442.04.

Applicant is notified that any subsequent amendment to the specification and/or claims must comply with 37 CFR 1.173(b).

Claim Rejections - 35 USC § 112

The following is a quotation of 35 U.S.C. 112(b):
(b) CONCLUSION.—The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the inventor or a joint inventor regards as the invention.

The following is a quotation of 35 U.S.C. 112 (pre-AIA), second paragraph:
The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claims 117-127 are rejected under 35 U.S.C. 112(b) or 35 U.S.C. 112 (pre-AIA), second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which the inventor or a joint inventor, or for pre-AIA the applicant regards as the invention.

Claim 117 recites the limitation "the GPU" in the last line. There is insufficient antecedent basis for this limitation in the claim.

Claim Rejections - 35 USC § 103

In the event the determination of the status of the application as subject to AIA 35 U.S.C. 102 and 103 (or as subject to pre-AIA 35 U.S.C. 102 and 103) is incorrect, any correction of the statutory basis for the rejection will not be considered a new ground of rejection if the prior art relied upon, and the rationale supporting the rejection, would be the same under either status.

The following is a quotation of pre-AIA 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 21-40, 56-60, 62-66, 137-138, and 140-141 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,477,256 B1 *(herein **Johnson**)* in further view of US 7,525,547 B1 *(herein **Diard**)*.

Claim 21

Buck discloses an accelerator system *(Buck: figure 2, element 190)* comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, the sequence of computations representing an artificial neural network, computations in the sequence of computations representing layers of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*;

a memory, operably coupled to the GPU, to store the input data and the output data *(Buck: figure 5, element 230)*; and

wherein each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory *(Buck: column 15, lines 46-62)*, the first texture holding a first value computed by the neuron during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture holding a second value computed by the neuron during the current time step of the sequence of computations *(Buck: column 16, lines 3-59) (Buck: generally column 13, line 23 to column 21, line 34)*.

Johnson demonstrates, to the extent ***Buck*** may not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the GPU and the memory, to transfer the input data into the memory and to transfer at least a portion of the output data from the memory" *(Johnson:*

figure 4A, elements 410 and 442; and column 7, lines 33-55). It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* may not explicitly state "to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the GPU", ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. The claim language defines "sequence of computations" as the whole sequence of artificial neural network computations and individual "computations" as respective layers of the neural network. Further, *Buck* demonstrates repeatedly performing the sequence of computations (of the whole artificial neural network) and periodically transferring input and output during those sequences of computations as a whole *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time; see also figure 5, learning module 545 training loop, and see "W" and "L" transferring in and out during the loop)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with output data transferred from the accelerator memory to the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application

of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 22

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the GPU is operably coupled to a central processing unit (CPU) of a computer system *(Buck: figure 2)*.

Claim 23

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the GPU is mounted on an expansion card operably coupled to a motherboard of a computer system *(Buck: figure 2, element 210; column 9, lines 15-22, AGP and PCI Express busses are for expansion cards coupled to a motherboard; Johnson: figure 3B)*.

Claim 24

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, further comprising: a central processing unit (CPU) *(Buck: figure 1, element 120)*, a system memory *(Buck: figure 1, element 130)*, and a non-volatile memory installed on a motherboard *(Buck: figure 1, element 133; column 6, lines 27-40, at least ISA bus indicates motherboard; Johnson: figure 3B)*.

Claim 25

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 24, wherein the GPU is operably coupled to the CPU via a system bus on the motherboard *(Buck: figure 1, element 190 containing the GPU; figure 2, elements 190 and 210; column 9, lines 15-22, AGP and PCI Express busses are for expansion cards coupled to a motherboard; Johnson: figure 3B)*.

Claim 26

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 25, wherein the system bus has a bandwidth of 3.2 GB/s to 15 GB/s *(Buck: column 9, lines 15-22, PCI Express bus has a bandwidth higher than 3.2GB/s; Johnson: column 3, lines 51-65)*.

Claim 27

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 25, wherein the controller is configured to transfer the input data and the at least a portion of the output data via the system bus *(Buck: figure 2, element 210)*.

Claim 28

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 25, further comprising: a video card, operably coupled to the motherboard via a local bus, to transfer graphical information to a display *(Buck: figure 1, elements 190, 121, and 191; claim language "local bus" is reference language "system bus")*.

Claim 29

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 28, wherein the local bus has a lower bandwidth than the system bus *(Buck: figure 1, element 121; column 6, lines 27-40, at least ISA bus; claim language "local bus" is reference language "system bus"; figure 2, element 210; column 9, lines 15-22; ISA bus has a lower bandwidth than a PCI Express bus)*.

Claim 30

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the memory comprises: a first memory partition to store parameters common to every computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4; see also figure 9 and column 15, line 46 to column 16, line 59)*.

Claim 31

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 30, wherein the memory further comprises: a second memory partition to store data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4; see also figure 9 and column 15, line 46 to column 16, line 59)*.

Claim 32

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 31, wherein the second memory partition comprises: a first partition to hold external input data patterns *(Buck: figure 5, element 530; column 10, lines 61-62)*; a second partition to hold representations of internal variables *(Buck: column 15, lines 49-60)*; a third partition to hold an input of the computation in the sequence of computations *(Buck: column 10, lines 43-44)*; and a fourth partition to hold the output of the computation in the sequence of computations *(Buck: figure 5, element 510)*.

Claim 33

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 31, wherein the second memory partition is configured to accumulate the output of the computation in the sequence of computations

during execution, by the GPU, of computational elements of the computation in the sequence of computations *(Buck: column 11, lines 50-52)*.

Claim 34

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 31, wherein the second memory partition is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to computational elements executed by the GPU during the computation in the sequence of computations *(Buck: figure 5, at least element 510; column 11, line 12 to column 12, line 15)*.

Claim 35

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the computation in the sequence of computations performed by the GPU comprises a plurality of computational elements *(Buck: column 21, lines 61-64)*.

Claim 36

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 35, wherein the plurality of computational elements represents a layer of neurons in an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27; output of a neural network is an output layer of neurons)*.

Claim 37

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 36, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 38

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the memory is configured to store intermediate results of the sequence of computations *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)* and the controller is configured to transfer the output data from the memory to a main memory without transferring the intermediate results *(Buck: figure 5, temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W)*.

Claim 39

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the controller is configured to initiate transfer of at least a portion of the input data *(Buck: column 10, lines 36-43)* and to transfer the at least a portion of the output data in parallel *(figure 5, element 585; and column 6, lines 35-40, at least ISA bus is 8-bit wide indicating at least a portion of data is in parallel)*.

Claim 40

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 21, wherein the input data is acquired by at least one of a video camera, a microphone, or a cell recording electrode *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64)*.

Claim 56

Buck discloses an accelerator system *(Buck: figure 2, element 190)* comprising:

a graphics processing unit (GPU) to perform a sequence of computations on input data so as to generate output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*,

a memory, operably coupled to the GPU, to store parameters common to all of the computations in the sequence of computations and data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4; see also figure 9 and column 15, line 46 to column 16, line 59)*

each computation in the sequence of computations representing a corresponding layer of neurons in an artificial neural network *(Buck: column 21, lines 61-64; Buck: figure 7; column 22, lines 26-27; output of a neural network is an output layer of neurons)*

Johnson demonstrates, to the extent ***Buck*** may not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the GPU and the memory, to control performance of the sequence of computations and data exchange by the GPU, to transfer at least a portion of the input data into the memory" and "to transfer at least a portion of the output data from the memory" *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* may not explicitly state to transfer at least a portion of the data between memory "while the GPU performs the sequence of computations", ***Diard*** demonstrates that it was known at the time of invention to transfer output from an accelerator memory to the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. The claim language defines individual "computations" in a "sequence of computations" as respective layers of the neural network. Further, *Buck* demonstrates repeatedly performing the sequence of computations (of the whole artificial neural network) and periodically transferring input and output during those sequences of computations as a whole *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time; see also figure 5, learning module 545 training loop, and see "W" and "L" transferring in and out during the loop)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with data transferred between the accelerator memory and the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 57

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein all neurons in the layer of neurons are described by the same equation *(Buck: figure 7, same neuron type divided by layer; column 13, lines 33-55; and column 11, lines 33-36)*.

Claim 58

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein the memory is configured to hold: external input data patterns *(Buck: figure 5, element 530; column 10, lines 61-62)*; representations of internal variables *(Buck: column 15, lines 49-60; see also figure 9 and column 15, line 46 to column 16, line 59)*; an input to the computation in the sequence of computations *(Buck: column 10, lines 43-44)*; and the output of the computation in the sequence of computations *(Buck: figure 5, element 510)*.

Claim 59

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein the memory is configured to accumulate the output of the computation in the sequence of computations during execution of a plurality of computational elements by the GPU *(Buck: column 11, lines 50-52; see also figure 9 and column 15, line 46 to column 16, line 59)*.

Claim 60

Buck, ***Johnson*** and ***Diard*** disclose the accelerator system of claim 56, wherein the memory is configured to store an output of a previous computation in the sequence of computations, the output of the previous computation being accessible to a plurality of computational elements during the computation in the sequence of computations *(Buck: figure 5, at least element 530; column 11, line 12 to column 12, line 15)*.

Claim 62

Buck discloses a method comprising:

receiving, at a central processing unit (CPU), input data acquired from an external system *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15; external systems ranging from input devices to the GPU)*;

initializing graphics processing unit (GPU) programs for performing a sequence of computations on the input data received by the CPU *(Buck: column 15, line column 14, line 17 to column 16, lines 60)*;

performing, with the GPU, a sequence of computations on input data to generate output data based on textures and shaders *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, computations in the sequence of computations representing respective layers of neurons in an artificial neural network *(Buck: column 21, lines 61-64; figure 7; column 22, lines 26-27)*,

controlling data exchange from the GPU to the CPU and from the CPU to the GPU upon initialization of the GPU programs *(Buck: column 10, line 18 to column 12, line 26)*;

storing, in the memory operably coupled to the GPU, parameters common to all of the computations in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*;

storing, in the memory, data specific to the computation in the sequence of computations *(Buck: figure 2, 230; figure 5, elements 530, 535, 595, 510; column 10, lines 36-63; see also column 20, lines 1-4)*.

Johnson demonstrates, to the extent ***Buck*** may not explicitly state, that it was known at the time of invention to provide a controller operably coupled to a GPU *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element

and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* may not explicitly state transferring input and output data between memories during performance of the sequence of computations by the GPU, ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 4, lines 5-9; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. The claim language defines individual "computations" in a "sequence of computations" as respective layers of the neural network. Further, *Buck* demonstrates repeatedly performing the sequence of computations (of the whole artificial neural network) and periodically transferring input and output during those sequences of computations as a whole *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time; see also figure 5, learning module 545 training loop, and see "W" and "L" transferring in and out during the loop)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and controller of *Johnson* with data transferred between accelerator memory and the main memory during performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 63-66

The limitations of claims 63-66 correspond to the limitations of claims 57-61 and as such are rejected in a corresponding manner.

Claim 137

Buck discloses a system for executing an artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*, the system comprising:

a first processing unit to receive input data in real time from a sensor *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 21, lines 54-64; and also column 11, line 12 to column 12, line 15)*;

a memory, operably coupled to the first processing unit, to store a first portion of the input data *(Buck: figure 5, element 230)*;

a second processing unit, operably coupled to the memory, to generate first output data from the first portion of the input data, the first output data representing an output of at least one neuron in a first layer of neurons in the artificial neural network *(Buck: figure 7; column 21, lines 61-64; column 22, lines 26-27)*

and allow a user to change the input data on the fly during the execution of the artificial neural network *(Buck: column 9, lines 33-36, "interactive use applications, such as speech recognition"; considering Buck, through multiple cycles, routinely exchanges information with the neural network of the GPU, and speech is interactive, that is changing on the fly)*.

Johnson demonstrates, to the extent ***Buck*** does not explicitly state, that it was known at the time of invention to provide a controller, "operably coupled to the second processing unit and the memory, to store a second portion of the input data in the memory" *(Johnson: figure 4A, elements 410 and 442; and column 7, lines 33-55)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement accelerator and GPU of *Buck* with a controller for transferring data as suggested by *Johnson*'s teachings. This implementation would have been obvious because one of ordinary skill in

the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both are directed to graphics processing systems.

To the extent *Buck* does not explicitly transferring data "while the second processing unit generates the first output data", ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations and periodically transferring input and output during those sequences of computations *(Buck: column 11, line 12 to column 12, line 26; figure 5, inputs and outputs; column 11, line 59 to column 12, line 6; column 23, lines 7-9; and column 10, lines 33-35, GPU and CPU compute at the same time; see also figure 5, learning module 545 training loop, and see "W" and "L" transferring in and out during the loop)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* and the controller of *Johnson* with data transferred between memories while generating data as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 138

Buck, ***Johnson*** and ***Diard*** disclose the system of claim 137, wherein the sensor is a microphone *(Buck: column 7, lines 51-59)*.

Claim 140

Buck, ***Johnson*** and ***Diard*** disclose the system of claim 137, wherein the second processing unit is further configured to generate second output data from the second portion of the input data, the second output data representing an output of at least one neuron in a second layer of neurons in the artificial neural network *(Buck: figure 5, inputs and outputs; column 11, line 59 to column 12, line 6; figure 7; column 21, line 65 to column 22, line 27)*.

Claim 141

Buck, ***Johnson*** and ***Diard*** disclose the system of claim 140, further comprising: a random access memory, operably coupled to the first processing unit, to receive the first output data from the memory while the second processing unit generates the second output data *(Buck: figure 1, element 132; column 6, lines 53-59; column 7, lines 7-18; column 11, line 59 to column 12, line 6; and column 10, lines 33-35, GPU and CPU compute at the same time)*.

Claims 90, 92-94, and 108-111 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,525,547 B1 *(herein **Diard**)*.

Claim 90

Buck discloses a method of performing a sequence of computations with a processing unit *(Buck: figure 1, elements 120 and 160; figure 2, element 120; column 9, lines 33-36; figure 1, elements 121, 130, 141, 152, 156; figure 1, element 190; figure 2, elements 220 and 230)*, the method comprising:

receiving input data *(Buck: figure 1, elements 120 and 160; figure 2, element 120; column 9, lines 33-36; figure 5; column 21, lines 54-64)*;

performing, by the processing unit, the sequence of computations on a first portion of the input data so as to generate a first portion of output data *(Buck: figure 2, element 220; figure 5; figure 7; column 12, lines 16-26)*, intermediate computations in the sequence of computations yielding intermediate results *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W)*, the first portion of the output data representing an output of an artificial neural network, wherein performing the sequence of computations on the first portion of the input data comprises accumulating a first value for an output variable in a first texture by a first computational element *(Buck: figure 7, output elements; column 16, lines 3-34; column 13, line 25 to column 14, line 44)*;

transferring a second portion of the input data from a first memory partition to a second memory partition *(Buck: column 11, line 12 to column 12, line 26)*;

transferring a second portion of the output data from the second memory partition to the first memory partition *(Buck: column 11, line 12 to column 12, line 26)*; and

performing, by the processing unit, the sequence of computations on the second portion of the input data so as to generate a third portion of output data *(Buck: column 12, lines 16-26)*, wherein performing the sequence of computations on the second portion of the input data comprises accumulating a second value for the output variable in the first texture by the first computational element and accessing the first value of the output variable in the first texture by a second computational element *(Buck: column 15, lines 57-60; column 16, lines 8-20; and column 16, lines 21-34)*.

To the extent *Buck* does not explicitly state "in parallel with performing the sequence of computations on the first portion of the input data by the processing unit, transferring a second portion of the input data from a first memory partition to a second memory partition" and "in parallel with performing the sequence of computations on the first portion of the input data by the processing unit, transferring a

second portion of the output data from the second memory partition to the first memory partition", ***Diard*** demonstrates that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing the sequence of computations (of the whole artificial neural network) and periodically transferring input and output during those sequences of computations as a whole *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time; see also figure 5, learning module 545 training loop, and see "W" and "L" transferring in and out during the loop)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* with data transferred between the accelerator memory and the main memory in parallel with performance of the sequence of computations as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 92

Buck and ***Diard*** disclose the method of claim 90, wherein receiving the input data comprises receiving the input data from at least one of a video camera, a microphone, or a cell recording electrode *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 23, lines 26-30)*.

Claim 93

Buck and ***Diard*** disclose the method of claim 90, wherein receiving the input data occurs at a first rate and performing the sequence of computations occurs at a second rate different than the first rate *(Buck:*

column 2, lines 55-67; column 9, lines 10-14; column 14, lines 40-44; CPU and GPU function at different rates).

Claim 94

Buck and ***Diard*** disclose the method of claim 90, wherein transferring the second portion of the output data comprises transferring the second portion of the output data from the second memory partition to the first memory partition without transferring any of the intermediate results from the second memory partition to the first memory partition *(Buck: figure 5, elements 595, temporary data, or 565, intermediate versions of W; figure 5, temporary data is never sent back; and column 11, lines 12-41, some embodiments do not sent back the intermediate versions of W).*

Claim 108

Buck discloses a method of performing a numerical simulation representing an artificial neural network *(Buck: column 3, lines 53-65; column 10, lines 43-53; column 11, lines 42-45; column 13, line 25 to column 14, line 15; column 21, lines 60-67; column 23, lines 10-14)*, the method comprising:

performing, by a processing unit, at least one computation during a first computational cycle of the numerical simulation *(Buck: column 11, lines 29-41; column 12, lines 16-26)*, the first computational cycle representing a first layer of neurons in the artificial neural network *(Buck: figure 7; column 22, lines 1-24)*;

storing, by a memory operably coupled to the processing unit, data used to perform the at least one computation *(Buck: figure 5, element 230)*;

controlling performance of the at least one computation by the processing unit *(Buck: figure 5, element 545)*;

controlling transfer of the output data from the processing unit to the memory during the first computational cycle *(Buck: figure 5, element 220 joined to 230)*;

the second computational cycle representing a second layer of neurons in the artificial neural network *(Buck: figure 7; column 22, lines 1-24)*.

Diard demonstrates, to the extent ***Buck*** does not explicitly state "while performing the at least one computation, acquiring input data with a central processing unit and operably coupled to a main memory" and "controlling transfer of input data to the memory from the main memory for use by the processing unit during a second computational cycle of the numerical simulation", that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing cycles of computation and periodically transferring input and output during cycles *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time; see also figure 5, learning module 545 training loop, and see "W" and "L" transferring in and out during the loop)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* with data transferred between the accelerator memory and the main memory while the GPU performs computations, calculations, and computational cycles as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 109

Buck discloses the method of claim 108, further comprising: dictating an order of execution of instructions to the processing unit *(Buck: column 11, lines 42-58; column 14, lines 9-66)*.

Claim 110

Diard demonstrates, to the extent ***Buck*** does not explicitly state "transferring output data from the memory to the main memory while the processing unit performs the at least one computation", that it was known at the time of invention to transfer data between an accelerator memory and the main memory while the GPU is performing other operations *(Diard: figure 1, transfer arrows; column 5, lines 1-3; column 6, lines 10-26; column 13, lines 4-25)*. Further, *Buck* demonstrates repeatedly performing cycles of computation and periodically transferring input and output during cycles *(Buck: column 11, line 12 to column 12, line 26; and column 10, lines 33-35, GPU and CPU compute at the same time)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement the CPU and GPU system of *Buck* with data transferred between the accelerator memory and the main memory while the GPU performs at least one computation as suggested by *Diard*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both *Buck* and *Diard* are directed to CPU and GPU processing.

Claim 111

Buck discloses the method of claim 108, further comprising: performing successive computational cycles of the numerical simulation by feeding output data generated by the processing unit from a previous computational cycle and input data for a next computational cycle into the processing unit *(Buck: column 11, lines 29-41; column 12, lines 16-26)*.

Claims 41 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,477,256 B1 *(herein **Johnson**)* in view of US 7,525,547 B1 *(herein **Diard**)* in further view of US 7,119,810 B2 *(herein **Sumanaweera**)*.

Claim 41

To the extent ***Buck*** does not explicitly state the accelerator system of claim 21, wherein the input data is loaded from disk. ***Sumanaweera*** demonstrates that it was known at the time of invention to load GPUs from disk *(Sumanaweera: column 1, lines 10-19)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement ***Buck*** with input data loaded from disk as suggested by ***Sumanaweera***'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art.

Claim 91 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,525,547 B1 *(herein **Diard**)* in further view of US 7,119,810 B2 *(herein **Sumanaweera**)*.

Claim 91

The limitations of claim 91 correspond to the limitations of claim 41 and as such are rejected in a corresponding manner.

Claim 139 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of US 7,477,256 B1 *(herein **Johnson**)* in view of US 7,525,547 B1 *(herein **Diard**)* in further view of US 2006/0129506 A1 *(herein **Edelman**)*.

Claim 139

Buck does not explicitly state wherein the sensor is a video camera. ***Edelman*** demonstrates that it was known at the time of invention to make use of video camera input sensors *(Edelman: paragraph 0049)* in an environment related to artificial neural networks *(Edelman: paragraph 0006)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement an input sensor of *Buck* with or as a video camera as suggested by *Edelman*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both references are directed to artificial neural networks.

Claim 147 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 7,219,085 B2 *(herein **Buck**)* in view of *(herein **Diard**)* in further view of US 2006/0129506 A1 *(herein **Edelman**)*.

Claim 147

Buck and ***Diard*** disclose the method of claim 108, wherein acquiring the input data comprises acquiring the input data in real time *(Buck: column 7, line 51 to column 8, line 2; column 9, lines 33-36; column 23, lines 26-30)*.

Buck does not explicitly state wherein the sensor is a video camera. ***Edelman*** demonstrates that it was known at the time of invention to make use of video camera input sensors *(Edelman: paragraph 0049)* in an environment related to artificial neural networks *(Edelman: paragraph 0006)*. It would have been obvious to one of ordinary skill in the art at the time of invention to implement an input sensor of *Buck* with or as a video camera as suggested by *Edelman*'s teachings. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a substitution and application of one known element and technique for another yielding a predictable result using an acceptable piece of prior art; and both references are directed to artificial neural networks.

Response to Arguments

Patent Owner's arguments filed 04/16/2020 *(herein Remarks)* have been fully considered but they are not persuasive. Claims 1, 12, 42, 95, 112, and 117 are allowed as indicated below. The Remarks directed toward claims 21, 56, 62, 90, 108, and 137 are addressed as follows.

With regard to claim 21, the cited prior art does disclose the limitations as presented. The Remarks indicate the cited prior art does not show transfers during performance of the sequence of computations. Unlike claim 1, the instant claim does not require a "to transfer … during performance of the intermediate computations in the sequence of computations by the at least one graphics processing unit" limitation, where "intermediate computations" are defined as "representing respective layers of the artificial neural network and yielding intermediate results". Transfers in claim 21 are "during performance of the sequence of computations by the GPU" and the claim defines "sequence of computations" as the entire artificial network. Therefore, the transfers may happen at any time in the larger artificial network "sequence of computations" which is a broad enough term to include other

calculations or activities), and need not occur during the layer computations. This argument is not persuasive.

With regard to claim 56, the cited prior art does disclose the limitations as presented for the same reasons as discussed for claim 21. Transfers in the instant claim may happen at any time in the larger artificial network "sequence of computations" (which is a broad enough term to include other calculations or activities not related to the layer computations), and need not occur during the layer computations. This argument is not persuasive.

With regard to claim 62, the cited prior art does disclose the limitations as presented for the same reasons as discussed for claims 21 and 56. Additionally, the Remarks state "Buck, Diard, and Johnson also fail to disclose 'controlling, by [a] controller, data exchange from the GPU to the CPU and from the CPU to the GPU upon initialization of the GPU programs,' in part because Diard and Johnson are silent on transferring data from the GPU to the CPU". However, both *Johnson* and *Diard* teach transferring and sharing data from the GPU memory to with the system/CPU memory *(Johnson: column 7, lines 52-55; Diard: column 4, lines 5-9)*. This argument is not persuasive.

With regard to claim 90, the Remarks indicate certain limitations are not addressed in the preceding rejections. However, these limitations are found and addressed in at least claim 12 and its rejection. Further, the rejection of claim 90 points the corresponding rejection of claim 12. This argument is not persuasive.

With regard to claim 108, the Remarks indicate certain limitations are not addressed in the preceding rejections, including: "acquire input data while calculating a first layer of a neural network with a processing unit" and "to transfer that input data from main memory to the processing unit memory for the processing unit to use when calculating a second layer of the neural network". However, *Buck* in view of *Diard* does show these limitations. As described above, *Buck* transfers data between CPU and GPU memory repeatedly through many cycles (loop iterations) of an artificial neural

network. In *Buck* each cycle can represent both a first layer and a second layer of the neural network. As shown above, *Diard* can transfer data while the GPU is going something else. The particular language of this claim is broad with regard to "computational cycle" and who/what is "controlling". Also, the claim doesn't actually mention transferring data, but instead "storing" data, which is a broader term. Additionally, the claimed "input data" acquired is broad, and reads upon any data the CPU processes in parallel with the GPU computations and then sends to the GPU. This argument is not persuasive.

With regard to claim 137, the cited prior art does disclose the limitations as presented for the same reasons as discussed for claims 21 and 56. The Response indicates the prior art "doesn't yield a computer system with controller that transfers output data from the GPU to the CPU while computing neural network layers" and "Buck discloses transferring neural network output data from the GPU to the CPU but is silent on making those transfers after the GPU computes a first layer and before the GPU computes a second layer in the neural network". However, these asserted limitations are not found in claim 137. The claim doesn't actually transfer anything. This argument is not persuasive.

Allowable Subject Matter

Claims 1-11, 12-14, 16-20, 42-55, 95-99, 112-114, 116, and 142-144 are allowed.

The following is a statement of reasons for the indication of allowable subject matter: the cited prior art of record does not reasonably disclose or suggest the combined limitations of the independent claims 1, 12, 42, 95, and 112. With respect to independent claim 1, ***Buck*** in view of ***Johnson*** and ***Diard*** does not clearly show the following limitations as arranged in the claim: "at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, the sequence of computations representing an artificial neural network, intermediate computations in the sequence of computations representing respective layers of the artificial neural network and yielding intermediate results"; and "a controller … to control performance

of the sequence of computations by the at least one graphics processing unit, to transfer the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the sequence of computations by the at least one graphics processing unit, and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one graphics processing unit".

With respect to independent claim 12, the cited prior art does not clearly show the following limitations as arranged in the claim: "performing, by the GPU, the sequence of computations on a first portion of input data so as to generate a first portion of the output data, the first portion of the output data representing an output of a neuron in a first layer of the artificial neural network …"; "in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory …"; "in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory …, the second portion of the output data representing an output of a neuron in a second layer in the artificial neural network"; and "making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements during the second time step".

With respect to independent claim 42, the cited prior art does not clearly show the following limitations as arranged in the claim: "performing, by the graphics processing unit (GPU), a first computation in the sequence of computations on the first input data … to generate first output data, computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of the first computation in the sequence of computations representing an output of a first neuron in a first layer in the artificial neural network"; and "transferring second input data acquired from the external system in real time into the memory operably coupled to the GPU after the GPU starts the first computation and before the GPU starts a second computation of the sequence

of computations, an output of the second computation in the sequence of computations representing an output of a second neuron in a second layer in the artificial neural network".

With respect to independent claim 95, the cited prior art does not clearly show the following limitations as arranged in the claim: "a camera to generate input data in real time"; "a processing unit, operably coupled to the camera … to perform the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network"; "to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit"; and "to transfer a second portion of the output data to the first memory partition in parallel with performance the sequence of computations by the processing unit". In claim 95, the input data is more than just trainable parameters and therefore *Buck* figure 5 does not apply.

With respect to independent claim 112, ***Buck*** in view of ***Johnson*** does not clearly show the following limitations as arranged in the claim: "a memory, operably coupled to the CPU, to store the first input data in a first partition, referenced by a first pointer, before computing a first layer of neurons of the artificial neural network"; "a processing unit … to perform, during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data, the first output data representing an output of at least one neuron in the first layer of neurons"; and "a controller … to … transfer the first output data to another memory during computation of the second layer of neurons".

Additionally, claims 117-127 would be allowable if rewritten or amended to overcome the rejection(s) under 35 U.S.C. 112(b) or 35 U.S.C. 112 (pre-AIA), 2nd paragraph, set forth in this Office action.

With respect to independent claim 117, the cited prior art does not clearly show the following limitations as arranged in the claim: "at least one processing unit to perform a sequence of computations representing an artificial neural network on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations representing layers of the neural network and yielding intermediate results"; and "a controller … to control transfer of the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the sequence of computations by the at least one processing unit, to control transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one processing unit, and to control performance of the sequence of computations by the GPU".

Conclusion

Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Correspondence Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to William H. Wood whose telephone number is (571)272-3736 and fax number is (571)273-3736. The examiner can normally be reached on Monday-Friday 7:30am-4pm. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Alexander Kosowski can be reached on (571)272-3744. The Central Reexamination Unit's (CRU) fax number is (571)273-9900.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR systems, see https://ppair.uspto.gov/epatent/portal/home. For questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/William H. Wood/
Primary Examiner, CRU 3992

Conferees:

/RSD/

/ALEXANDER J KOSOWSKI/
Supervisory Patent Examiner, Art Unit 3992

Applicant-Initiated Interview Summary	Application No. 15/808,201	Applicant(s) Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	AIA (FITF) Status No

All participants (applicant, applicants representative, PTO personnel):

(1) WILLIAM H. WOOD.

(2) Max Colice.

(3) Alexander Kosowski.

(4) Rachna Desai.

Date of Interview: 04 March 2020.

Type: ☐ Telephonic ☑ Video Conference
☐ Personal [copy given to: ☐ applicant ☐ applicant's representative]

Exhibit shown or demonstration conducted: ☑ Yes ☐ No.
If Yes, brief description: Agenda.

Issues Discussed ☐101 ☐112 ☐102 ☑103 ☐Others
(For each of the checked box(es) above, please describe below the issue and detailed description of the discussion)

Claim(s) discussed: 1,68,108 and 137.

Identification of prior art discussed: Buck.

Substance of Interview
(For each issue discussed, provide a detailed description and indicate if agreement was reached. Some topics may include: identification or clarification of a reference or a portion thereof, claim interpretation, proposed amendments, arguments of any applied references etc...)

No agreement was reached, but progress made in understanding the general breadth of the claim language and possible views regarding memory transfers between the GPU and CPU, related to: defining "sequence of computations" and other broad language; and whether Buck would allow a significant number of memory transfers and at the right time. Additional discussion involved: how to view "intermediate results" and "output data"; interpretations of "streams"; and artificial intelligence layers and stages of operation. Ultimately, it is not yet clear how some of the above terminology must be viewed in relation to the cited prior art and Patent Owner's amendments, arguments, and/or exhibits might be persuasive once fully presented..

Applicant recordation instructions: The formal written reply to the last Office action must include the substance of the interview. (See MPEP section 713.04). If a reply to the last Office action has already been filed, applicant is given a non-extendable period of the longer of one month or thirty days from this interview date, or the mailing date of this interview summary form, whichever is later, to file a statement of the substance of the interview.

Examiner recordation instructions: Examiners must summarize the substance of any interview of record. A complete and proper recordation of the substance of an interview should include the items listed in MPEP 713.04 for complete and proper recordation including the identification of the general thrust of each argument or issue discussed, a general indication of any other pertinent matters discussed regarding patentability and the general results or outcome of the interview, to include an indication as to whether or not agreement was reached on the issues raised.

☑ Attachment

/William H. Wood/ Primary Examiner, Art Unit 3992	

Summary of Record of Interview Requirements

Manual of Patent Examining Procedure (MPEP), Section 713.04, Substance of Interview Must be Made of Record
A complete written statement as to the substance of any face-to-face, video conference, or telephone interview with regard to an application must be made of record in the application whether or not an agreement with the examiner was reached at the interview.

Title 37 Code of Federal Regulations (CFR) 1.133 Interviews
Paragraph (b)

In every instance where reconsideration is requested in view of an interview with an examiner, a complete written statement of the reasons presented at the interview as warranting favorable action must be filed by the applicant. An interview does not remove the necessity for reply to Office action as specified in §§ 1.111, 1.135. (35 U.S.C. 132)

37 CFR §1.2 Business to be transacted in writing.
All business with the Patent or Trademark Office should be transacted in writing. The personal attendance of applicants or their attorneys or agents at the Patent and Trademark Office is unnecessary. The action of the Patent and Trademark Office will be based exclusively on the written record in the Office. No attention will be paid to any alleged oral promise, stipulation, or understanding in relation to which there is disagreement or doubt.

The action of the Patent and Trademark Office cannot be based exclusively on the written record in the Office if that record is itself incomplete through the failure to record the substance of interviews.

It is the responsibility of the applicant or the attorney or agent to make the substance of an interview of record in the application file, unless the examiner indicates he or she will do so. It is the examiners responsibility to see that such a record is made and to correct material inaccuracies which bear directly on the question of patentability.

Examiners must complete an Interview Summary Form for each interview held where a matter of substance has been discussed during the interview by checking the appropriate boxes and filling in the blanks. Discussions regarding only procedural matters, directed solely to restriction requirements for which interview recordation is otherwise provided for in Section 812.01 of the Manual of Patent Examining Procedure, or pointing out typographical errors or unreadable script in Office actions or the like, are excluded from the interview recordation procedures below. Where the substance of an interview is completely recorded in an Examiners Amendment, no separate Interview Summary Record is required.

The Interview Summary Form shall be given an appropriate Paper No., placed in the right hand portion of the file, and listed on the "Contents" section of the file wrapper. In a personal interview, a duplicate of the Form is given to the applicant (or attorney or agent) at the conclusion of the interview. In the case of a telephone or video-conference interview, the copy is mailed to the applicants correspondence address either with or prior to the next official communication. If additional correspondence from the examiner is not likely before an allowance or if other circumstances dictate, the Form should be mailed promptly after the interview rather than with the next official communication.

The Form provides for recordation of the following information:
- Application Number (Series Code and Serial Number)
- Name of applicant
- Name of examiner
- Date of interview
- Type of interview (telephonic, video-conference, or personal)
- Name of participant(s) (applicant, attorney or agent, examiner, other PTO personnel, etc.)
- An indication whether or not an exhibit was shown or a demonstration conducted
- An identification of the specific prior art discussed
- An indication whether an agreement was reached and if so, a description of the general nature of the agreement (may be by attachment of a copy of amendments or claims agreed as being allowable). Note: Agreement as to allowability is tentative and does not restrict further action by the examiner to the contrary.
- The signature of the examiner who conducted the interview (if Form is not an attachment to a signed Office action)

It is desirable that the examiner orally remind the applicant of his or her obligation to record the substance of the interview of each case. It should be noted, however, that the Interview Summary Form will not normally be considered a complete and proper recordation of the interview unless it includes, or is supplemented by the applicant or the examiner to include, all of the applicable items required below concerning the substance of the interview.

A complete and proper recordation of the substance of any interview should include at least the following applicable items:
1) A brief description of the nature of any exhibit shown or any demonstration conducted,-
2) an identification of the claims discussed,
3) an identification of the specific prior art discussed,
4) an identification of the principal proposed amendments of a substantive nature discussed, unless these are already described on the Interview Summary Form completed by the Examiner,
5) a brief identification of the general thrust of the principal arguments presented to the examiner,
(The identification of arguments need not be lengthy or elaborate. A verbatim or highly detailed description of the arguments is not required. The identification of the arguments is sufficient if the general nature or thrust of the principal arguments made to the examiner can be understood in the context of the application file. Of course, the applicant may desire to emphasize and fully describe those arguments which he or she feels were or might be persuasive to the examiner.)
6) a general indication of any other pertinent matters discussed, and
7) if appropriate, the general results or outcome of the interview unless already described in the Interview Summary Form completed by the examiner.

Examiners are expected to carefully review the applicants record of the substance of an interview. If the record is not complete and accurate, the examiner will give the applicant an extendable one month time period to correct the record.

Examiner to Check for Accuracy

If the claims are allowable for other reasons of record, the examiner should send a letter setting forth the examiners version of the statement attributed to him or her. If the record is complete and accurate, the examiner should place the indication, Interview Record OK on the paper recording the substance of the interview along with the date and the examiners initials.

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020					
	1	✓	✓	✓	=					
	2	✓	✓	✓	=					
	3	✓	✓	✓	=					
	4	✓	✓	✓	=					
	5	✓	✓	✓	=					
	6	✓	✓	✓	=					
	7	✓	✓	✓	=					
	8	✓	✓	✓	=					
	9	✓	✓	✓	=					
	10	✓	✓	✓	=					
	11	✓	✓	✓	=					
	12	✓	✓	✓	=					
	13	✓	✓	✓	=					
	14	✓	✓	✓	=					
	15	✓	✓	✓	-					
	16	✓	✓	✓	=					
	17	✓	✓	✓	=					
	18	✓	✓	✓	=					
	19	✓	✓	✓	=					
	20	✓	✓	✓	=					
	21	✓	✓	✓	✓					
	22	✓	✓	✓	✓					
	23	✓	✓	✓	✓					
	24	✓	✓	✓	✓					
	25	✓	✓	✓	✓					
	26	✓	✓	✓	✓					
	27	✓	✓	✓	✓					
	28	✓	✓	✓	✓					
	29	✓	✓	✓	✓					
	30	✓	✓	✓	✓					
	31	✓	✓	✓	✓					
	32	✓	✓	✓	✓					
	33	✓	✓	✓	✓					
	34	✓	✓	✓	✓					
	35	✓	✓	✓	✓					
	36	✓	✓	✓	✓					
	37	✓	✓	✓	✓					
	38	✓	✓	✓	✓					
	39	✓	✓	✓	✓					
	40	✓	✓	✓	✓					
	41	✓	✓	✓	✓					
	42	✓	✓	✓	=					

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020					
	43	✓	✓	✓	=					
	44	✓	✓	✓	=					
	45	✓	✓	✓	=					
	46	✓	✓	✓	=					
	47	✓	✓	✓	=					
	48	✓	✓	✓	=					
	49	✓	✓	✓	=					
	50	✓	✓	✓	=					
	51	✓	✓	✓	=					
	52	✓	✓	✓	=					
	53	✓	✓	✓	=					
	54	✓	✓	✓	=					
	55	✓	✓	✓	=					
	56	✓	✓	✓	✓					
	57	✓	✓	✓	✓					
	58	✓	✓	✓	✓					
	59	✓	✓	✓	✓					
	60	✓	✓	✓	✓					
	61	✓	-	-	-					
	62	✓	✓	✓	✓					
	63	✓	✓	✓	✓					
	64	✓	✓	✓	✓					
	65	✓	✓	✓	✓					
	66	✓	✓	✓	✓					
	67	✓	-	-	-					
	68	✓	✓	✓	-					
	69	✓	✓	✓	-					
	70	✓	✓	✓	-					
	71	✓	✓	✓	-					
	72	✓	✓	✓	-					
	73	✓	✓	✓	-					
	74	✓	✓	✓	-					
	75	✓	✓	✓	-					
	76	✓	✓	✓	-					
	77	✓	✓	✓	-					
	78	✓	✓	✓	-					
	79	✓	✓	✓	-					
	80	✓	✓	✓	-					
	81	✓	✓	✓	-					
	82	✓	✓	✓	-					
	83	✓	✓	✓	-					
	84	✓	✓	✓	-					
	85	✓	✓	✓	-					
	86	✓	✓	✓	-					
	87	✓	✓	✓	-					
	88	✓	✓	✓	-					
	89	✓	✓	✓	-					
	90	✓	✓	✓	✓					
	91	✓	✓	✓	✓					
	92	✓	✓	✓	✓					
	93	✓	✓	✓	✓					
	94	✓	✓	✓	✓					

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020					
	95	✓	✓	✓	=					
	96	✓	✓	✓	=					
	97	✓	✓	✓	=					
	98	✓	✓	✓	=					
	99	✓	✓	✓	=					
	100	✓	✓	✓	-					
	101	✓	✓	✓	-					
	102	✓	✓	✓	-					
	103	✓	✓	✓	-					
	104	✓	✓	✓	-					
	105	✓	✓	✓	-					
	106	✓	✓	✓	-					
	107	✓	✓	✓	-					
	108	✓	✓	✓	✓					
	109	✓	✓	✓	✓					
	110	✓	✓	✓	✓					
	111	✓	✓	✓	✓					
	112	✓	✓	✓	=					
	113	✓	✓	✓	=					
	114	✓	✓	✓	=					
	115	✓	✓	✓	-					
	116	✓	✓	✓	=					
	117	✓	✓	✓	✓					
	118	✓	✓	✓	✓					
	119	✓	✓	✓	✓					
	120	✓	✓	✓	✓					
	121	✓	✓	✓	✓					
	122	✓	✓	✓	✓					
	123	✓	✓	✓	✓					
	124	✓	✓	✓	✓					
	125	✓	✓	✓	✓					
	126	✓	✓	✓	✓					
	127	✓	✓	✓	✓					
	128	✓	✓	✓	-					
	129	✓	✓	✓	-					
	130	✓	✓	✓	-					
	131	✓	✓	✓	-					
	132	✓	✓	✓	-					
	133	✓	✓	✓	-					
	134	✓	✓	✓	-					
	135	✓	✓	✓	-					
	136	✓	✓	✓	-					
	137		✓	✓	✓					
	138		✓	✓	✓					
	139		✓	✓	✓					
	140		✓	✓	✓					
	141		✓	✓	✓					
	142			✓	=					
	143			✓	=					
	144			✓	=					
	145			✓	-					
	146			✓	-					

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020					
	147			✓	✓					
	148			✓	-					

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

CPC - Searched*		
Symbol	Date	Examiner
G06F 9/5027	10/10/2018	/WHW/
G06T 1/20	10/10/2018	/WHW/
G06T 1/60	10/10/2018	/WHW/
G06N 99/005	10/10/2018	/WHW/
G06N 3/063	10/10/2018	/WHW/
G06F 2209/509	10/10/2018	/WHW/

CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

/William H. Wood/ Primary Examiner, Art Unit 3992	

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

Search Notes		
Search Notes	**Date**	**Examiner**
Prosecution History Review	10/10/2018	/WHW/
IDS Review	10/10/2018	/WHW/
Litigation Search	10/10/2018	/WHW/
EAST BRS Keyword Search	10/10/2018	/WHW/
NPL IEEE Search	10/10/2018	/WHW/
IDS Review	05/03/2019	/WHW/
EAST BRS	05/03/2019	/WHW/
EAST BRS	12/10/2019	/WHW/
IDS Review	12/10/2019	/WHW/
EAST BRS update	05/13/2020	/WHW/
NPL update	05/13/2020	/WHW/

Interference Search			
US Class/CPC Symbol	**US Subclass/CPC Group**	**Date**	**Examiner**

/William H. Wood/ Primary Examiner, Art Unit 3992	

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S167	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2020/04/28 15:18
S168	1	S167 and parallel	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/28 15:18
S169	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2020/04/29 15:31
S170	1	S169 and layer	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/29 15:31
S171	1	S169 and intermediate	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/29 15:57
S172	1	("9189828").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2020/05/01 16:13
S173	1	S172 and input	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/01 16:13
S174	13309	(graphics adj processing adj unit) and accelerator	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:48

S176	33780	(multip$4 or plural$4 or several or many) near3 (memory adj (bank or partition or section or module))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:50
S177	57846	(graphics adj processing adj unit) with memory	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:50
S178	6049	(S174 or S176 or S177) and ((transfer$4 or moving or shifting or copying or send$4 or swap$4) near3 (during or while or parallel))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:54
S179	0	S178 and (layer near3 (arificial adj neural adj network))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:55
S180	0	S178 and (layer with (arificial adj neural adj network))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:55
S181	178	S178 and (layer with neuron)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:55
S182	176	S178 and (layer near3 neuron)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:55
S183	174	(S181 or S182) and ((layer or neural or neuron) with (while or during or parallel))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:57

S184	13	(S181 or S182) and ((layer or neural or neuron) with (simultaneous$4 or concurrent$4))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:57
S185	175	(S181 or S182) and ((layer or neural or neuron) same (while or during or parallel))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 13:59
S186	169	S182 and ((input or output) near3 (during or while or simultaneous$4 or concurrent$4 or parallel))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/05/07 14:01

EAST Search History (Interference)

<This search history is empty>

5/13/2020 1:19:55 PM
C:\Users\wwood\Documents\EAST\Workspaces\15 808201.wsp

Wood, William

From: Max Colice <colice@smithbaluch.com>
Sent: Monday, March 2, 2020 12:45 PM
To: Wood, William
Subject: US 15/808,201 - Attorney Docket No. NRLA-001US03 - Proposed Interview Agenda

Examiner Wood,

Thank you again for granting a WebEx interview in US 15/808,201 on Wednesday, March 4, at 10 am ET. During the interview, I would like to speak to you and your colleagues about three points:

(1) Using a controller to control data transfer from a GPU to a CPU while the GPU executes the neural network. This data transfer may include transferring output data from one layer of neurons in the neural network *from accelerator memory to main memory* while the GPU is executing the neural network. Here are some of the ways we have tried to capture this transfer in the claims:

Claim 1: "a controller … to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the sequence of computations by the at least one graphics processing unit." (See also method claim 42.)

Claim 56: "a controller … to transfer at least a portion of the output data from the memory in parallel with performance of the sequence of computations by the GPU." (See also claims 62, 90, and 128.)

Claim 108: "performing, by a processing unit, at least one computation during a first computational cycle of the numerical simulation, ***the first computational cycle representing a first layer of neurons in the artificial neural network***;
…
controlling transfer of the output data from the processing unit to the memory during the first computational cycle." (See also claim 117.)

These claims have been rejected as obvious over Buck in view of (Johnson and) Diard, with Diard allegedly demonstrating that it was well known to transfer data between accelerator memory and main memory while the GPU was performing other operations.

Buck is silent on transferring output data from one layer of neurons in the neural network *from accelerator memory to main memory while the GPU is executing the neural network*. In fact, Buck says that "data transfer from GPU to CPU is quite expensive" (13:15), which arguably teaches away from transferring data from the GPU memory to main memory while the GPU is performing calculations.

Diard is also silent on transferring output data *from accelerator memory to main memory* while the GPU is executing calculations. In Diard, the GPU renders images from data provided by the CPU. The data transfer is all one way, from the CPU to the GPU, and never from the GPU to the CPU, at least not while the GPU is performing calculations. The GPU output is rendered on a display after the calculations are done. Diard's FIG. 1 shows that this display is coupled directly to the GPUs, so there's no reason for the GPU to transfer data back to the CPU.

(2) Shifting control of data exchange, etc., between user interaction and computational streams while executing a neural network as in claims 68 and 79. These claims have been rejected as obvious over Buck in view of Johnson and Diard, with Diard again providing data transfer ***between*** accelerator and main memory.

(3) Changing input data "on-the-fly" as in claim 137. According to the Office Action, Buck's brief mention of "speech recognition" qualifies as "changing input data on the fly during execution of the artificial neural network."

Recognizing that Internet communications are not secure, I hereby authorize the USPTO to communicate with the undersigned and practitioners in accordance with 37 CFR 1.33 and 37 CFR 1.34 concerning any subject matter of this application by video conferencing, instant messaging, or electronic mail. I understand that a copy of these communications will be made of record in the application file.

Max Colice
Reg. No. 65,634

NVIDIA Ex1002 Page 1001 of 1081

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 15/808,201	Filing Date 11/09/2017	☐ To be Mailed

ENTITY: ☐ LARGE ☑ SMALL ☐ MICRO

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A		N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A		N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*		x $40 =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*		x $210 =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).				
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))					
* If the difference in column 1 is less than zero, enter "0" in column 2.				TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT	09/03/2020	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 56	Minus	** 139	= 0		x $50 =	0
	Independent (37 CFR 1.16(h))	* 6	Minus	*** 16	= 0		x $230 =	0
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=		x $0 =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=		x $0 =	
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.	LIE
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".	/ANNETTE COWAN/
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".	
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.	

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchechnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT AFTER FINAL

In response to the Final Office Action dated June 5, 2020, Applicant submits the following Amendments and Remarks, and respectfully requests reconsideration of the application in view thereof.

Amendments to the Claims are reflected in the listing of the claims which begins on page 2 of this paper.

Remarks begin on page 17 of this paper.

Electronic Acknowledgement Receipt

EFS ID:	40462363
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice/Lindsey Lovato
Filer Authorized By:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03
Receipt Date:	03-SEP-2020
Filing Date:	09-NOV-2017
Time Stamp:	13:46:49
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		NRLA-001US03_Amendment_After_Final_due_2020905.pdf	305600 50882fbbe4f29a6d7b4231e69562c4673874c07c	yes	19

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Applicant Arguments/Remarks Made in an Amendment	17	19
	Claims	2	16
	Response After Final Action	1	1
Warnings:			
Information:			
	Total Files Size (in bytes):	305600	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

REMARKS

In response to the Final Office Action dated June 5, 2020, in the above-identified application, Applicant respectfully requests reconsideration. All of the issues raised in the Office Action have been carefully considered and are addressed in this response. The application as now presented is believed to be in allowable condition.

1. Status of the Claims and Summary of Amendments

Upon entry of this response, claims 1–14, 15–20, 42–55, 95–99, 112–114, 116–127, and 142–144 will be pending in this application, of which claims 1, 12, 42, 95, 112, and 117 are independent claims.

Claim 117 is being amended to correct an antecedent basis error.

Claims 21–41, 56–60, 62–66, 90–94, 108–111, 137–141, and 147 are being canceled solely to expedite allowance and without prejudice or disclaimer. Applicant reserves the right to pursue the subject matter of these claims in a continuing application.

2. Prior and Concurrent Judicial Proceedings

To Applicant's knowledge, there are no prior or concurrent judicial proceedings involving U.S. Patent No. 9,189,828.

3. Rejections Under 35 U.S.C. § 112

Claims 117–127 were rejected as indefinite because there is insufficient antecedent basis for the term "the GPU" in independent claim 117. As indicated above, claim 117 is being amended to recite "the at least one processing unit" instead of "the GPU," overcoming this rejection.

4. Allowed Claims

Applicant thanks Examiner Wood for indicating that claims 1–14, 42–55, 95–99, 112–114, 116, and 142–147 are allowed and that claims 117–127 would be allowed if amended to overcome the indefiniteness rejection.

5. Rejections Under 35 U.S.C. § 103

Claims 21–40, 56–60, 62–66, 137, 138, 140, and 141 have been rejected as obvious over US 7,219,085 B2 to Buck in view of US 7,477,256 B1 to Johnson and US 7,525,547 B1 to Diard.

Claims 90, 92–94, and 108–111 have been rejected as unpatentable over Buck in view of Diard.

Claim 41 has been rejected as unpatentable over Buck in view of Johnson, Diard, and US 7,119,810 to Sumanaweera.

Claim 91 has been rejected as unpatentable over Buck in view of Diard and Sumanaweera.

Claim 139 has been rejected as unpatentable over Buck in view of Johnson, Diard, and US 2006/0129506 A1 to Edelman.

Claim 147 has been rejected as unpatentable over Buck in view of Diard and Edelman.

Applicant respectfully traverses these rejections for at least the reasons given in Applicant's last response. Nevertheless, solely to expedite allowance, Applicant is canceling these claims to obviate these rejections.

CONCLUSION

It is respectfully believed that all of the rejections set forth in the Office Action have been addressed. However, the absence of a reply to a specific rejection, objection, or comment set forth in the Office Action does not signify agreement with or concession of that rejection, objection, or comment. In addition, there may be reasons for patentability of any or all pending claims (or other claims) that have not been expressed. Furthermore, nothing in this paper should be construed as an intent to concede any issue with regard to any claim.

Applicant believes that the present application is now in condition for allowance. Favorable reconsideration of the application as amended is respectfully requested.

The Examiner is invited to contact the undersigned by telephone if it is felt that a telephone interview would advance the prosecution of the present application.

The Director is hereby authorized to charge any appropriate fees under 37 C.F.R. §§ 1.16, 1.17, and 1.21 that may be required by this paper and to credit any overpayment, to Deposit Account No. 60-1876. If an extension of time is required to make this reply timely, please charge

the appropriate extension of time fee under 37 C.F.R. § 1.17 to Deposit Account No. 60-1876 if not already paid via EFS.

Dated: September 3, 2020

Respectfully submitted,

Smith Baluch LLP
376 Boylston St, Suite 401
Boston MA 02116

Phone: (617) 947-7280

By /Christopher Max Colice/
Christopher Max Colice
Registration No. 65,634

IN THE CLAIMS:

Set forth below in ascending order, with status identifiers, is a complete listing of all claims currently under examination. Changes to any amended claims are indicated by [[double brackets]], ~~strikethrough~~, and/or underlining. This listing also reflects any cancellation and/or addition of claims.

1. (Four Times Amended) A computer system, comprising:

a central processing unit to receive input data;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the [first] main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, the sequence of computations representing an artificial neural network, intermediate computations in the sequence of computations representing respective layers of the artificial neural network and yielding intermediate results; and

accelerator memory, operably coupled to the [graphic] at least one graphics processing unit, to store the results of the [plurality of sequential] sequence of computations; and

a controller, operably coupled to the at least one graphics processing unit and the accelerator memory, to initialize textures and shaders in the accelerator memory for performing the sequence of computations, to control performance of the sequence of computations by the at least one graphics processing unit, to transfer the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the sequence of computations by the at least one graphics processing unit, and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one [graphic] graphics processing unit.

2. The computer system of claim 1, wherein the central processing unit is configured to receive the input data in response to a user interaction.

3. The computer system of claim 1, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one graphics processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

4. The computer system of claim 1, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

5. The computer system of claim 1, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

6. The computer system of claim 1, wherein accelerator memory comprises:

a first memory bank to store parameters common to all of the computations in the sequence of computations; and

a second memory bank to store data specific to at least one computation in the sequence of computations.

7. The computer system of claim 1, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

8. The computer system of claim 1, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one graphics processing unit has begun to perform another sequence of computations.

9. The computer system of claim 8, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one graphics processing unit.

10. The computer system of claim 1, wherein the controller is configured to control execution of the sequence of computations by the at least one graphics processing unit.

11. (Amended) The computer system of claim 1, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central [processor] processing unit, to acquire the input data in real time.

12. (Thrice Amended) A method of performing a sequence of computations representing an artificial neural network on a computer system comprising a central processing unit (CPU), a main memory operably coupled to the central processing unit via a bus, an accelerator operably coupled to the CPU and the main memory via the bus, the accelerator comprising a graphics processing unit (GPU) and an accelerator memory, the method comprising:

(A) performing, by the GPU, the sequence of computations on a first portion of [the] input data so as to generate a first portion of [the] output data, the first portion of the output data representing an output of a neuron in a first layer of the artificial neural network, intermediate computations in the sequence of computations yielding intermediate results, wherein performing the sequence of computations on the first portion of the input data comprises (i) assigning an output variable to a first texture and a second texture, the output variable being included in a first computational element of a plurality of computational elements, the plurality of computational elements representing the sequence of computations and (ii) accumulating a first value for the output variable in the first texture during a first time step;

(B) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory to the accelerator via the bus; [and]

(C) in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory to the main memory via the bus, the second portion of the output data representing an output of a neuron in a second layer in the artificial neural network; and

(D) performing, by the GPU, the sequence of computations on the second portion of the input data, wherein performing the sequence of computations on the second portion of the input data comprises (i) accumulating a second value for the output variable in the second texture during a second time step and (ii) making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements during the second time step.

13. The method of claim 12, further comprising:

storing the input data in the main memory in response to a user interaction.

14. The method of claim 12, further comprising:

receiving the input data at a first rate; and

wherein (A) comprises performing the sequence of computations at a second rate different than the first rate.

15. Cancel claim 15.

16. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory without transferring any of the intermediate results of the plurality of sequential computations from the accelerator memory to the main memory so as to reduce data transfer via the bus.

17. The method of claim 12, wherein (C) comprises:

transferring the second portion of the output data from the accelerator memory to the main memory after the GPU has begun to perform another sequence of computations.

18. The method of claim 17, wherein (C) further comprises:

initiating transfer of the second portion of the output data in parallel with performance of at least one computation in the other sequence of computations.

19. The method of claim 12, further comprising:

acquiring the input data in real time with at least one of a video camera, a microphone, or a cell recording electrode operably coupled to the CPU.

20. The method of claim 12, further comprising:

storing parameters common to all of the computations in the sequence of computations in a first memory bank in the accelerator memory; and

storing data specific to at least one computation in the sequence of computations in a second memory bank in the accelerator memory.

21. Cancel claim 21.

22. Cancel claim 22.

23. Cancel claim 23.

24. Cancel claim 24.

25. Cancel claim 25.

26. Cancel claim 26.

27. Cancel claim 27.

28. Cancel claim 28.

29. Cancel claim 29.

30. Cancel claim 30.

31. Cancel claim 31.

32. Cancel claim 32.

33. Cancel claim 33.

34. Cancel claim 34.

35. Cancel claim 35.

36. Cancel claim 36.

37. Cancel claim 37.

38. Cancel claim 38.

39. Cancel claim 39.

40. Cancel claim 40.

41. Cancel claim 41.

42. (New, Thrice Amended) A method of performing a sequence of computations representing an artificial neural network, the method comprising:

receiving, at a central processing unit (CPU), first input data acquired from an external system in real time;

initializing, by a controller operably coupled to a graphics processing unit (GPU), textures and shaders in a memory operably coupled to the GPU;

transferring the first input data received by the CPU to the memory operably coupled to the GPU;

performing, by the graphics processing unit (GPU), a first computation in the sequence of computations on the first input data based on the textures and shaders to generate first output data, computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of the first computation in the sequence of computations representing an output of a first neuron in a first layer in the artificial neural network;

storing, in the memory operably coupled to the GPU, the first input data and the first output data; and

transferring second input data acquired from the external system in real time into the memory operably coupled to the GPU after the GPU starts the first computation and before the GPU starts a second computation of the sequence of computations, an output of the second computation in the sequence of computations representing an output of a second neuron in a second layer in the artificial neural network.

43. (New, Amended) The method of claim 42, wherein transferring the second input data comprises transferring the second input data via a bus operably coupled to the CPU.

44. (New, Amended) The method of claim 42, further comprising:

transferring the first output data from the memory to another memory during the second computation in the sequence of computations.

45. (New) The method of claim 44, further comprising:

storing intermediate results of the sequence of computations in the memory, and

wherein transferring the first output data from the memory to the other memory occurs without transferring the intermediate results of the sequence of computations.

46. (New) The method of claim 44, wherein transferring the second input data and transferring the first output data occurs in parallel.

47. (New) The method of claim 42, further comprising:

storing, in a first memory partition of the memory, parameters common to all of the computations in the sequence of computations.

48. (New) The method of claim 47, further comprising:

storing, in a second memory partition of the memory, data specific to the first computation in the sequence of computations.

49. (New) The method of claim 48, further comprising:

storing, in the second memory partition, external input data patterns, representations of internal variables, an input of the computation in the sequence of computations, and the output of the computation in the sequence of computations.

50. (New) The method of claim 42, wherein storing the first output data comprises:

accumulating, in the memory, outputs of computational elements executed by the GPU in performing the first computation in the sequence of computations.

51. (New) The method of claim 42, further comprising:

storing, in the memory, an output of a previous computation in the sequence of computations; and

accessing, by the GPU, the output of the previous computation during performance of the computation in the sequence of computations.

52. (New) The method of claim 42, wherein performing the first computation comprises executing a plurality of computational elements representing a layer of neurons in an artificial neural network.

53. (New) The method of claim 52, wherein all neurons in the layer of neurons are described by the same equation.

54. (New, Amended) The method of claim 42, further comprising:

acquiring the second input data with at least one of a video camera, a microphone, or a cell recording electrode.

55. (New, Twice Amended) The method of claim 42, further comprising:

loading the second input data from disk.

56. Cancel claim 56.

57. Cancel claim 57.

58. Cancel claim 58.

59. Cancel claim 59.

60. Cancel claim 60.

61. Cancel claim 61.

62. Cancel claim 62.

63. Cancel claim 63.

64. Cancel claim 64.

65. Cancel claim 65.

66. Cancel claim 66.

67. Cancel claim 67.

68. Cancel claim 68.

69. Cancel claim 69.

70. Cancel claim 70.

71. Cancel claim 71.

72. Cancel claim 72.

73. Cancel claim 73.

74. Cancel claim 74.

75. Cancel claim 75.

76. Cancel claim 76.

77. Cancel claim 77.

78. Cancel claim 78.

79. Cancel claim 79.

80. Cancel claim 80.

81. Cancel claim 81.

82. Cancel claim 82.

83. Cancel claim 83.

84. Cancel claim 84.

85. Cancel claim 85.

86. Cancel claim 86.

87. Cancel claim 87.

88. Cancel claim 88.

89. Cancel claim 89.

90. Cancel claim 90.

91. Cancel claim 91.

92. Cancel claim 92.

93. Cancel claim 93.

94. Cancel claim 94.

95. (New, Thrice Amended) A system for performing a sequence of computations, the system comprising:

a camera to generate input data in real time;

a first memory partition;

a second memory partition operably coupled to the first memory partition; and

a processing unit, operably coupled to the camera, the first memory partition, and the second memory partition, to perform the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network,

wherein the first memory partition is configured to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit,

wherein the second memory partition is configured to transfer a second portion of the output data to the first memory partition in parallel with performance the sequence of computations by the processing unit, and

wherein the sequence of computations represents the artificial neural network, each neuron in the artificial neural network has an output variable assigned to a first texture and a second texture in the memory, the first texture holds a first value of the output variable computed during a previous time step of the sequence of computations and accessible to other neurons in the neural network during a current time step of the sequence of computations and the second texture accumulates a second value of the output variable computed during the current time step.

96. (New) The system of claim 95, wherein the first memory partition and the second memory partition are logical partitions.

97. (New) The system of claim 95, wherein the processing unit is comprises a graphics processing unit (GPU).

98. (New) The system of claim 95, wherein the processing unit is configured to receive the input data at a first rate and to perform the sequence of computations at a second rate is different than the first rate.

99. (New) The system of claim 95, wherein the second memory partition is configured to transfer the second portion of the output data to the first memory partition without transferring any of the intermediate results to the first memory partition.

100. Cancel claim 100.

101. Cancel claim 101.

102. Cancel claim 102.

103. Cancel claim 103.

104. Cancel claim 104.

105. Cancel claim 105.

106. Cancel claim 106.

107. Cancel claim 107.

108. Cancel claim 108.

109. Cancel claim 109.

110. Cancel claim 110.

111. Cancel claim 111.

112. (New, Thrice Amended) A system for executing an artificial neural network, the system comprising:

a central processing unit (CPU) to provide first input data;

a memory, operably coupled to the CPU, to store the first input data in a first partition, referenced by a first pointer, before computing a first layer of neurons of the artificial neural network;

a processing unit, operably coupled to the memory, to perform, during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data, the first output data representing an output of at least one neuron in the first layer of neurons; and

a controller, operably coupled to the processing unit and the memory, to:

store the first output data in a second partition of the memory, the second partition referenced by a second pointer, and to swap the first pointer with the second pointer at the end of the computation of the first layer of neurons, such that the first output data becomes an input for a second layer of neurons of the artificial neural network,

transfer the first output data to another memory during computation of the second layer of neurons, and

dictate an order of execution of instructions to the processing unit to perform the computation of the first layer of neurons.

113. (New) The system of claim 112, wherein the processing unit comprises a graphics processing unit.

114. (New) The system of claim 112, wherein the controller is configured to send instructions for performing the at least one calculation to the processing unit.

115. Cancel claim 115.

116. (New, Amended) The system of claim 112, wherein the memory further comprises:

a third partition to store internal variables; and

a fourth partition to store data used as input at a particular layer of neurons of the artificial neural network.

117. (New, Four Times Amended) A computer system, comprising:

a central processing unit to receive input data acquired from an external system;

main memory, operably coupled to the central processing unit via a bus, to store the input data received by the central processing unit;

an accelerator, operably coupled to the central processing unit and the main memory via the bus, to receive at least a portion of the input data from the main memory, the accelerator comprising:

at least one processing unit to perform a sequence of computations representing an artificial neural network on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations representing layers of the neural network and yielding intermediate results; and

accelerator memory, operably coupled to the at least one processing unit, to store the results of the sequence of computations; and

a controller, operably coupled to the at least one processing unit and the accelerator memory, to control transfer of the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the

sequence of computations by the at least one processing unit, to control transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one processing unit, and to control performance of the sequence of computations by the at least one processing unit.

118. (New) The computer system of claim 117, wherein the central processing unit is configured to receive the input data in response to a user interaction.

119. (New) The computer system of claim 117, wherein:

the central processing unit is configured to receive the input data at a first rate; and

the at least one processing unit is configured to perform the sequence of computations at a second rate different than the first rate.

120. (New) The computer system of claim 117, wherein the main memory is configured to store a copy of the output data stored in the accelerator memory.

121. (New) The computer system of claim 117, wherein an output of at least one computation in the sequence of computations represents an output of at least one neuron in an artificial neural network.

122. (New) The computer system of claim 117, wherein accelerator memory comprises:

a first memory partition to store parameters common to all of the computations in the sequence of computations; and

a second memory partition to store data specific to at least one computation in the sequence of computations.

123. (New) The computer system of claim 117, wherein the controller is configured to transfer the output data from the accelerator memory to the main memory without transferring any of the intermediate results from the accelerator memory to the main memory so as to reduce data transfer via the bus.

124. (New) The computer system of claim 117, wherein the controller is configured to transfer at least a portion of the output data from the accelerator memory to the main memory after the at least one processing unit has begun to perform another sequence of computations.

125. (New) The computer system of claim 124, wherein the controller is configured to initiate transfer of the at least a portion of the input data and to transfer the at least a portion of the output data in parallel with performance of at least one computation in the other sequence of computations by the at least one processing unit.

126. (New) The computer system of claim 117, wherein the controller is configured to control execution of the sequence of computations by the at least one processing unit.

127. (New, Amended) The computer system of claim 117, further comprising:

at least one of a video camera, a microphone, or a cell recording electrode, operably coupled to the central processing unit, to acquire the input data in real time.

128. Cancel claim 128.

129. Cancel claim 129.

130. Cancel claim 130.

131. Cancel claim 131.

132. Cancel claim 132.

133. Cancel claim 133.

134. Cancel claim 134.

135. Cancel claim 135.

136. Cancel claim 136.

137. Cancel claim 137.

138. Cancel claim 138.

139. Cancel claim 139.

140. Cancel claim 140.

141. Cancel claim 141.

142. (New) The computer system of claim 1, wherein the controller is configured to inform the central processing unit that the sequence of computations is finished.

143. (New) The computer system of claim 1, wherein the controller is configured to reduce a processing load on the central processing unit.

144. (New) The computer system of claim 1, wherein the controller is configured to reduce interactions between the central processing unit and the accelerator.

145. Cancel claim 145.

146. Cancel claim 146.

147. Cancel claim 147.

148. Cancel claim 148.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 09/15/2020

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

NOTIFICATION DATE	DELIVERY MODE
09/15/2020	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docketing@smithbaluch.com
eofficeaction@appcoll.com
smith@smithbaluch.com

Advisory Action Before the Filing of an Appeal Brief	**Application No.** 15/808,201	**Applicant(s)** Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	**Art Unit** 3992	**AIA (FITF) Status** No

--The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

THE REPLY FILED 03 September 2020 FAILS TO PLACE THIS APPLICATION IN CONDITION FOR ALLOWANCE.

NO NOTICE OF APPEAL FILED

1. ☑ The reply was filed after a final rejection. No Notice of Appeal has been filed. To avoid abandonment of this application, applicant must timely file one of the following replies: (1) an amendment, affidavit, or other evidence, which places the application in condition for allowance; (2) a Notice of Appeal (with appeal fee) in compliance with 37 CFR 41.31; or (3) a Request for Continued Examination (RCE) in compliance with 37 CFR 1.114 if this is a utility or plant application. Note that RCEs are not permitted in design applications. The reply must be filed within one of the following time periods:
   a) ☐ The period for reply expires ____ months from the mailing date of the final rejection.
   b) ☑ The period for reply expires on: (1) the mailing date of this Advisory Action; or (2) the date set forth in the final rejection, whichever is later. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of the final rejection.
   c) ☐ A prior Advisory Action was mailed more than 3 months after the mailing date of the final rejection in response to a first after-final reply filed within 2 months of the mailing date of the final rejection.The current period for reply expires ____ months from the mailing date of *the prior Advisory Action* or SIX MONTHS from the mailing date of the final rejection, whichever is earlier.
   Examiner Note: If box 1 is checked, check either box (a), (b) or (c). ONLY CHECK BOX (b) WHEN THIS ADVISORY ACTION IS THE FIRST RESPONSE TO APPLICANTS FIRST AFTER-FINAL REPLY WHICH WAS FILED WITHIN TWO MONTHS OF THE FINAL REJECTION. ONLY CHECK BOX (c) IN THE LIMITED SITUATION SET FORTH UNDER BOX (c). See MPEP 706.07(f).

Extensions of time may be obtained under 37 CFR 1.136(a). The date on which the petition under 37 CFR 1.136(a) and the appropriate extension fee have been filed is the date for purposes of determining the period of extension and the corresponding amount of the fee. The appropriate extension fee under 37 CFR 1.17(a) is calculated from: (1) the expiration date of the shortened statutory period for reply originally set in the final Office action; or (2) as set forth in (b) or (c) above, if checked. Any reply received by the Office later than three months after the mailing date of the final rejection, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

NOTICE OF APPEAL

2. ☐ The Notice of Appeal was filed on _____. A brief in compliance with 37 CFR 41.37 must be filed within two months of the date of filing the Notice of Appeal (37 CFR 41.37(a)), or any extension thereof (37 CFR 41.37(e)), to avoid dismissal of the appeal. Since a Notice of Appeal has been filed, any reply must be filed within the time period set forth in 37CFR 41.37(a).

AMENDMENTS

3. ☐ The proposed amendments filed after a final rejection, but prior to the date of filing a brief, will not be entered because
   a) ☐ They raise new issues that would require further consideration and/or search (see NOTE below);
   b) ☐ They raise the issue of new matter (see NOTE below);
   c) ☐ They are not deemed to place the application in better form for appeal by materially reducing or simplifying the issues for appeal; and/or
   d) ☐ They present additional claims without canceling a corresponding number of finally rejected claims.
   NOTE: _____ (See 37CFR 1.116 and 41.33(a)).
4. ☐ The amendments are not in compliance with 37CFR 1.121. See attached Notice of Non-Compliant Amendment (PTOL-324).
5. ☐ Applicants reply has overcome the following rejection(s): _____
6. ☐ Newly proposed or amended claim(s) _____ would be allowable if submitted in a separate, timely filed amendment canceling the non-allowable claim(s).
7. ☑ For purposes of appeal, the proposed amendment(s):(a) ☐ will not be entered, or (b) ☑ will be entered, and an explanation of how the new or amended claims would be rejected is provided below or appended.

AFFIDAVIT OR OTHER EVIDENCE

8. ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____
9. ☐ The affidavit or other evidence filed after final action, but before or on the date of filing a Notice of Appeal will not be entered because applicant failed to provide a showing of good and sufficient reasons why the affidavit or other evidence is necessary and was not earlier presented. See 37 CFR 1.116(e).
10. ☐ The affidavit or other evidence filed after the date of filing the Notice of Appeal, but prior to the date of filing a brief, will not be entered because the affidavit or other evidence failed to overcome all rejections under appeal and/or appellant fails to provide a showing of good and sufficient reasons why it is necessary and was not earlier presented. See 37 CFR 41.33(d)(1).
11. ☐ The affidavit or other evidence is entered. An explanation of the status of the claims after entry is below or attached.

REQUEST FOR RECONSIDERATION/OTHER

12. ☑ The request for reconsideration has been considered but does NOT place the application in condition for allowance because: See Continuation Sheet.
13. ☐ Note the attached Information *Disclosure Statement*(s). (PTO/SB/08) Paper No(s). _____
14. ☐ Other: _____.

STATUS OF CLAIMS

15. The status of the claim(s) is (or will be) as follows:
    Claim(s) allowed:_____.
    Claim(s) objected to:1-14,16-20,42-55,95-99,112-114,116-127 and 142-144.
    Claim(s) rejected:_____.
    Claim(s) withdrawn from consideration:_____.

/William H. Wood/ Primary Examiner, Art Unit 3992	/A.J.K/ Supervisory Patent Examiner, Art Unit 3992

Continuation of REQUEST FOR RECONSIDERATION/OTHER 12. The request for reconsideration has been considered but does NOT place the application in condition for allowance because: The application is otherwise in condition for allowance, except the claims are objected to because the Terminal Disclaimer directed to 8,648,867 and filed 11/09/2017 is not approved. Appropriate correction (see Terminal Disclaimer review decision of 04/21/2020) is required before the application can finally progress to allowance.

Conferees:

/RSD/

OK TO ENTER: /WHW/

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchechnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT AFTER FINAL

In response to the Final Office Action dated June 5, 2020, Applicant submits the following Amendments and Remarks, and respectfully requests reconsideration of the application in view thereof.

Amendments to the Claims are reflected in the listing of the claims which begins on page 2 of this paper.

Remarks begin on page 17 of this paper.

OK TO ENTER: /WHW/

Attorney Docket No. NRLA-001US03 **PATENT**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re Application of: Anatoli Gorchetchnikov	Confirmation No.: 1077
Application No.: 15/808,201	Group Art Unit: 3992
Filed: 2017-11-09	Examiner: William H WOOD

FOR: **GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD**

VIA EFS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

REPLY TO ADVISORY ACTION AND TERMINAL DISCLAIMER

In response to the Advisory Action dated September 15, 2020 and the Terminal Disclaimer Review Decision dated April 21, 2020, Applicant submits the enclosed terminal disclaimer with respect to U.S. Patent No. 8,648,867. Applicant respectfully submits that this terminal disclaimer correctly identifies this application as the application being examined. Accordingly, Applicant respectfully asks for approval of this terminal disclaimer and allowance of this application.

No fee is believed due. If a fee is due, the Director is hereby requested and authorized to charge any fee to Deposit Account No. 60-1876, from which the undersigned is authorized to draw.

Dated: September 17, 2020

Respectfully submitted,
SMITH BALUCH LLP

SMITH BALUCH LLP
376 Boylston St., Suite 401
Boston, MA 02116

Tel: (617) 947-7280

By: /Christopher Max Colice/
Christopher Max Colice
Reg. No. 65,634

PTO/AIA/26 (04-14)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TERMINAL DISCLAIMER TO OBVIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT	Docket Number (Optional) NRLA-001US03

In re Application of: Anatoli Gorchetchnikov; Heather Marie Ames; Massimiliano Versace; Fabrizio Santi

Application No.: 15/808,201

Filed: 2017-11-09

For: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

The applicant, Neurala, Inc., owner of 100 percent interest in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of **prior patent** No. 8,648,867 as the term of said **prior patent** is presently shortened by any terminal disclaimer. The applicant hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the **prior patent** are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.

In making the above disclaimer, the applicant does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the **prior patent**, "as the term of said **prior patent** is presently shortened by any terminal disclaimer," in the event that said **prior patent** later:

expires for failure to pay a maintenance fee;
is held unenforceable;
is found invalid by a court of competent jurisdiction;
is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;
has all claims canceled by a reexamination certificate;
is reissued; or
is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer.

Check either box 1 or 2 below, if appropriate.

1. ☐ The undersigned is the applicant. If the applicant is an assignee, the undersigned is authorized to act on behalf of the assignee.

I hereby acknowledge that any willful false statements made are punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

2. ☒ The undersigned is an attorney or agent of record. Reg. No. 65634

/Christopher Max Colice/ — Signature
2020-09-17 — Date

Christopher Max Colice — Typed or printed name

Title

(617) 947-7280 — Telephone Number

☑ Terminal disclaimer fee under 37 CFR 1.20(d) included.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

This collection of information is required by 37 CFR 1.321. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	40582767
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice/Lindsey Lovato
Filer Authorized By:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03
Receipt Date:	17-SEP-2020
Filing Date:	09-NOV-2017
Time Stamp:	09:36:12
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Response After Final Action	NRLA-001US03_Terminal_Disclaimer_Transmittal_20200917.pdf	110478 b7f14770e0c306267aa5d5a20be36151023ad574	no	1

Warnings:

Information:					
2	Terminal Disclaimer Filed	NRLA-001US03_USPTO_AIA26_20200917_090523.pdf	155728 fc316b71429c24d139f8b68221c2e53226a67995	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	266206	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 15/808,201	Filing Date 11/09/2017	☐ To be Mailed

ENTITY: ☐ LARGE ☑ SMALL ☐ MICRO

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA		RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A		N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A		N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*		x $40 =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*		x $210 =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).				
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))					
* If the difference in column 1 is less than zero, enter "0" in column 2.				TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT	09/17/2020	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 56	Minus	** 139	= 0		x $50 =	0
	Independent (37 CFR 1.16(h))	* 6	Minus	*** 16	= 0		x $230 =	0
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)			
AMENDMENT		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA		RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=		x $0 =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=		x $0 =	
	☐ Application Size Fee (37 CFR 1.16(s))							
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							
							TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.	LIE
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".	/PATRICIA F LEWIS/
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".	
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.	

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Application Number	Application/Control No.	Applicant(s)/Patent under Reexamination	
* 15/808,201 *	15/808,201	Gorchetchnikov et al.	
	Examiner	**Art Unit**	
	WOOD, WILLIAM H	3992	

Document Code - DISQ	**Internal Document - DO NOT MAIL**

TERMINAL DISCLAIMER	☑ APPROVED	☐ DISAPPROVED
Date Filed: **17 September 2020**	**This patent is subject to a Terminal Disclaimer**	

Approved/Disapproved by:

/PAMELA A YOUNG/

Technology Center: PLRC

Telephone: (571)272-3622

_____

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

NOTICE OF ALLOWANCE AND FEE(S) DUE

149890 7590 09/30/2020

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER
WOOD, WILLIAM H

ART UNIT	PAPER NUMBER
3992	

DATE MAILED: 09/30/2020

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

TITLE OF INVENTION: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$600	$0.00	$0.00	$600	12/30/2020

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. PROSECUTION ON THE MERITS IS CLOSED. THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED. SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE DOES NOT REFLECT A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE IN THIS APPLICATION. IF AN ISSUE FEE HAS PREVIOUSLY BEEN PAID IN THIS APPLICATION (AS SHOWN ABOVE), THE RETURN OF PART B OF THIS FORM WILL BE CONSIDERED A REQUEST TO REAPPLY THE PREVIOUSLY PAID ISSUE FEE TOWARD THE ISSUE FEE NOW DUE.

HOW TO REPLY TO THIS NOTICE:

I. Review the ENTITY STATUS shown above. If the ENTITY STATUS is shown as SMALL or MICRO, verify whether entitlement to that entity status still applies.

If the ENTITY STATUS is the same as shown above, pay the TOTAL FEE(S) DUE shown above.

If the ENTITY STATUS is changed from that shown above, on PART B - FEE(S) TRANSMITTAL, complete section number 5 titled "Change in Entity Status (from status indicated above)".

For purposes of this notice, small entity fees are 1/2 the amount of undiscounted fees, and micro entity fees are 1/2 the amount of small entity fees.

II. PART B - FEE(S) TRANSMITTAL, or its equivalent, must be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted. If an equivalent of Part B is filed, a request to reapply a previously paid issue fee must be clearly made, and delays in processing may occur due to the difficulty in recognizing the paper as an equivalent of Part B.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Maintenance fees are due in utility patents issuing on applications filed on or after Dec. 12, 1980. It is patentee's responsibility to ensure timely payment of maintenance fees when due. More information is available at www.uspto.gov/PatentMaintenanceFees.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), by mail or fax, or via EFS-Web.

By mail, send to: Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

By fax, send to: (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

149890 7590 09/30/2020
SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being transmitted to the USPTO via EFS-Web or by facsimile to (571) 273-2885, on the date below.

(Typed or printed name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

TITLE OF INVENTION: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$600	$0.00	$0.00	$600	12/30/2020

EXAMINER	ART UNIT	CLASS-SUBCLASS
WOOD, WILLIAM H	3992	345-531000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

❑ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.

❑ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-09 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list
(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,
(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1 _____
2 _____
3 _____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document must have been previously recorded, or filed for recordation, as set forth in 37 CFR 3.11 and 37 CFR 3.81(a). Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE (B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent) : ❑ Individual ❑ Corporation or other private group entity ❑ Government

4a. Fees submitted: ❑ Issue Fee ❑ Publication Fee (if required) ❑ Advance Order - # of Copies _____

4b. Method of Payment: *(Please first reapply any previously paid fee shown above)*

❑ Electronic Payment via EFS-Web ❑ Enclosed check ❑ Non-electronic payment by credit card (Attach form PTO-2038)

❑ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment to Deposit Account No. _____

5. **Change in Entity Status** (from status indicated above)

❑ Applicant certifying micro entity status. See 37 CFR 1.29

❑ Applicant asserting small entity status. See 37 CFR 1.27

❑ Applicant changing to regular undiscounted fee status.

NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.
NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.
NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature _____ Date _____

Typed or printed name _____ Registration No. _____

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 09/30/2020

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

DATE MAILED: 09/30/2020

Determination of Patent Term Extension or Adjustment under 35 U.S.C. 154 (b)

A reissue patent is for "the unexpired part of the term of the original patent." See 35 U.S.C. 251. Accordingly, the above-identified reissue application is not eligible for Patent Term Extension or Adjustment under 35 U.S.C. 154(b).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

OMB Clearance and PRA Burden Statement for PTOL-85 Part B

The Paperwork Reduction Act (PRA) of 1995 requires Federal agencies to obtain Office of Management and Budget approval before requesting most types of information from the public. When OMB approves an agency request to collect information from the public, OMB (i) provides a valid OMB Control Number and expiration date for the agency to display on the instrument that will be used to collect the information and (ii) requires the agency to inform the public about the OMB Control Number's legal significance in accordance with 5 CFR 1320.5(b).

The information collected by PTOL-85 Part B is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450. Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Notice of Allowability	Application No. 15/808,201	Applicant(s) Gorchetchnikov et al.	
	Examiner WILLIAM H WOOD	Art Unit 3992	AIA (FITF) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☑ This communication is responsive to 09/03/2020.
   ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

3. ☑ The allowed claim(s) is/are 1-14,16-20,42-55,95-99,112-114,116-127 and 142-144 . As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see **http://www.uspto.gov/patents/init_events/pph/index.jsp** or send an inquiry to **PPHfeedback@uspto.gov.**

4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

   Certified copies:

   a) ☐ All b) ☐ Some *c) ☐ None of the:

   1. ☐ Certified copies of the priority documents have been received.
   2. ☐ Certified copies of the priority documents have been received in Application No. _____ .
   3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

   * Certified copies not received: _____ .

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____ .

   Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☐ Notice of References Cited (PTO-892)
2. ☐ Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date _____.
3. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material _____.
4. ☐ Interview Summary (PTO-413), Paper No./Mail Date. _____.
5. ☐ Examiner's Amendment/Comment
6. ☑ Examiner's Statement of Reasons for Allowance
7. ☐ Other _____.

/William H. Wood/
Primary Examiner, Art Unit 3992

The present application is being examined under the pre-AIA first to invent provisions.

ALLOWANCE

For reissue applications filed before September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the law and rules in effect on September 15, 2012. Where specifically designated, these are "pre-AIA" provisions.

For reissue applications filed on or after September 16, 2012, all references to 35 U.S.C. 251 and 37 CFR 1.172, 1.175, and 3.73 are to the current provisions.

Reissue

Applicant is reminded of the continuing obligation under 37 CFR 1.178(b), to timely apprise the Office of any prior or concurrent proceeding in which Patent No. 9,189,828 is or was involved. These proceedings would include interferences, reissues, reexaminations, and litigation.

Applicant is further reminded of the continuing obligation under 37 CFR 1.56, to timely apprise the Office of any information which is material to patentability of the claims under consideration in this reissue application.

These obligations rest with each individual associated with the filing and prosecution of this application for reissue. See also MPEP §§ 1404, 1442.01 and 1442.04.

Applicant is notified that any subsequent amendment to the specification and/or claims must comply with 37 CFR 1.173(b).

Allowable Subject Matter

Claims 1-11, 12-14, 16-20, 42-55, 95-99, 112-114, 116, 117-127, and 142-144 are allowed.

The following is a statement of reasons for the indication of allowable subject matter: the cited prior art of record does not reasonably disclose or suggest the combined limitations of the independent

claims 1, 12, 42, 95, and 112. With respect to independent claim 1, ***Buck*** in view of ***Johnson*** and ***Diard*** does not clearly show the following limitations as arranged in the claim: "at least one graphics processing unit to perform a sequence of computations on the at least a portion of the input data so as to generate output data, the sequence of computations representing an artificial neural network, intermediate computations in the sequence of computations representing respective layers of the artificial neural network and yielding intermediate results"; and "a controller … to control performance of the sequence of computations by the at least one graphics processing unit, to transfer the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the sequence of computations by the at least one graphics processing unit, and to transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one graphics processing unit".

With respect to independent claim 12, the cited prior art does not clearly show the following limitations as arranged in the claim: "performing, by the GPU, the sequence of computations on a first portion of input data so as to generate a first portion of the output data, the first portion of the output data representing an output of a neuron in a first layer of the artificial neural network …"; "in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the input data from the main memory …"; "in parallel with performing the sequence of computations by the GPU in (A), transferring a second portion of the output data from the accelerator memory …, the second portion of the output data representing an output of a neuron in a second layer in the artificial neural network"; and "making the first value of the output variable in the first texture accessible to other computational elements in the plurality of computational elements during the second time step".

With respect to independent claim 42, the cited prior art does not clearly show the following limitations as arranged in the claim: "performing, by the graphics processing unit (GPU), a first

computation in the sequence of computations on the first input data … to generate first output data, computations in the sequence of computations representing respective layers of neurons in the artificial neural network, an output of the first computation in the sequence of computations representing an output of a first neuron in a first layer in the artificial neural network"; and "transferring second input data acquired from the external system in real time into the memory operably coupled to the GPU after the GPU starts the first computation and before the GPU starts a second computation of the sequence of computations, an output of the second computation in the sequence of computations representing an output of a second neuron in a second layer in the artificial neural network".

With respect to independent claim 95, the cited prior art does not clearly show the following limitations as arranged in the claim: "a camera to generate input data in real time"; "a processing unit, operably coupled to the camera … to perform the sequence of computations on a first portion of the input data so as to generate a first portion of output data, intermediate computations in the sequence of computations yielding intermediate results, the first portion of the output data representing an output of an artificial neural network"; "to transfer a second portion of the input data to the second memory partition in parallel with performance the sequence of computations by the processing unit"; and "to transfer a second portion of the output data to the first memory partition in parallel with performance the sequence of computations by the processing unit". In claim 95, the input data is more than just trainable parameters and therefore *Buck* figure 5 does not apply.

With respect to independent claim 112, ***Buck*** in view of ***Johnson*** does not clearly show the following limitations as arranged in the claim: "a memory, operably coupled to the CPU, to store the first input data in a first partition, referenced by a first pointer, before computing a first layer of neurons of the artificial neural network"; "a processing unit … to perform, during computation of the first layer of neurons, at least one calculation on the first input data so as to generate first output data, the first output data representing an output of at least one neuron in the first layer of neurons"; and "a

controller … to … transfer the first output data to another memory during computation of the second layer of neurons".

With respect to independent claim 117, the cited prior art does not clearly show the following limitations as arranged in the claim: "at least one processing unit to perform a sequence of computations representing an artificial neural network on the at least a portion of the input data so as to generate output data, intermediate computations in the sequence of computations representing layers of the neural network and yielding intermediate results"; and "a controller … to control transfer of the at least a portion of the input data into the accelerator memory during performance of the intermediate computations in the sequence of computations by the at least one processing unit, to control transfer at least a portion of the output data from the accelerator memory to the main memory during performance of the intermediate computations in the sequence of computations by the at least one processing unit, and to control performance of the sequence of computations by the GPU".

Correspondence Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to William H. Wood whose telephone number is (571)272-3736 and fax number is (571)273-3736. The examiner can normally be reached on Monday-Friday 7:30am-4pm. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Alexander Kosowski can be reached on (571)272-3744. The Central Reexamination Unit's (CRU) fax number is (571)273-9900.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR systems, see https://ppair.uspto.gov/epatent/portal/home. For questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/William H. Wood/
Primary Examiner, CRU 3992

Conferees:

/RSD/

/ALEXANDER J KOSOWSKI/
Supervisory Patent Examiner, Art Unit 3992

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CPC						
Symbol					**Type**	**Version**
G06T	/	1	/	20	F	2013-01-01
G06F	/	9	/	5027	I	2013-01-01
G06T	/	1	/	60	I	2013-01-01
G06N	/	3	/	063	A	2013-01-01
G06N	/	20	/	00	A	2019-01-01
G06F	/	2209	/	509	A	2013-01-01

CPC Combination Sets							
Symbol					**Type**	**Set**	**Ranking**	**Version**
	/		/					

NONE		**Total Claims Allowed:**	
(Assistant Examiner)	(Date)	58	
/William H. Wood/ Primary Examiner, Art Unit 3992	22 September 2020	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	4

U.S. Patent and Trademark Office — Part of Paper No.: 20200909

NVIDIA Ex1002 Page 1045 of 1081

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

INTERNATIONAL CLASSIFICATION				
CLAIMED				
G06T	/	1	/	60
	/		/	

NON-CLAIMED				
G06F	/	9	/	50
G06T	/	1	/	20

US ORIGINAL CLASSIFICATION	
CLASS	SUBCLASS
1	1

CROSS REFERENCES(S)						
CLASS	SUBCLASS (ONE SUBCLASS PER BLOCK)					

NONE (Assistant Examiner)	(Date)	Total Claims Allowed: 58	
/William H. Wood/ Primary Examiner, Art Unit 3992 (Primary Examiner)	22 September 2020 (Date)	O.G. Print Claim(s) 1	O.G. Print Figure 4

U.S. Patent and Trademark Office

Part of Paper No.: 20200909

NVIDIA Ex1002 Page 1046 of 1081

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☑ T.D. ☐ R.1.47

CLAIMS

Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original
1	1	20	20		39		58		77	37	96		115		134
2	2		21		40		59		78	38	97	44	116		135
3	3		22		41		60		79	39	98	45	117		136
4	4		23	21	42		61		80	40	99	46	118		137
5	5		24	22	43		62		81		100	47	119		138
6	6		25	23	44		63		82		101	48	120		139
7	7		26	25	45		64		83		102	49	121		140
8	8		27	26	46		65		84		103	50	122		141
9	9		28	27	47		66		85		104	51	123	56	142
10	10		29	28	48		67		86		105	52	124	57	143
11	11		30	29	49		68		87		106	53	125	58	144
12	12		31	30	50		69		88		107	54	126		145
13	13		32	31	51		70		89		108	55	127		146
14	14		33	32	52		71		90		109		128		147
15	15		34	33	53		72		91		110		129		148
16	16		35	34	54		73		92		111		130		
17	17		36	35	55		74		93	41	112		131		
18	18		37		56		75		94	42	113		132		
19	19		38		57		76	36	95	43	114		133		

NONE		**Total Claims Allowed:**	
(Assistant Examiner)	(Date)	58	
/William H. Wood/ Primary Examiner, Art Unit 3992	22 September 2020	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	4

U.S. Patent and Trademark Office

Part of Paper No.: 20200909

NVIDIA Ex1002 Page 1047 of 1081

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020	09/22/2020				
	1	✓	✓	✓	=	=				
	2	✓	✓	✓	=	=				
	3	✓	✓	✓	=	=				
	4	✓	✓	✓	=	=				
	5	✓	✓	✓	=	=				
	6	✓	✓	✓	=	=				
	7	✓	✓	✓	=	=				
	8	✓	✓	✓	=	=				
	9	✓	✓	✓	=	=				
	10	✓	✓	✓	=	=				
	11	✓	✓	✓	=	=				
	12	✓	✓	✓	=	=				
	13	✓	✓	✓	=	=				
	14	✓	✓	✓	=	=				
	15	✓	✓	✓	-	-				
	16	✓	✓	✓	=	=				
	17	✓	✓	✓	=	=				
	18	✓	✓	✓	=	=				
	19	✓	✓	✓	=	=				
	20	✓	✓	✓	=	=				
	21	✓	✓	✓	✓	-				
	22	✓	✓	✓	✓	-				
	23	✓	✓	✓	✓	-				
	24	✓	✓	✓	✓	-				
	25	✓	✓	✓	✓	-				
	26	✓	✓	✓	✓	-				
	27	✓	✓	✓	✓	-				
	28	✓	✓	✓	✓	-				
	29	✓	✓	✓	✓	-				
	30	✓	✓	✓	✓	-				
	31	✓	✓	✓	✓	-				
	32	✓	✓	✓	✓	-				
	33	✓	✓	✓	✓	-				
	34	✓	✓	✓	✓	-				
	35	✓	✓	✓	✓	-				
	36	✓	✓	✓	✓	-				
	37	✓	✓	✓	✓	-				
	38	✓	✓	✓	✓	-				
	39	✓	✓	✓	✓	-				
	40	✓	✓	✓	✓	-				
	41	✓	✓	✓	✓	-				
	42	✓	✓	✓	=	=				

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020	09/22/2020				
	43	✓	✓	✓	=	=				
	44	✓	✓	✓	=	=				
	45	✓	✓	✓	=	=				
	46	✓	✓	✓	=	=				
	47	✓	✓	✓	=	=				
	48	✓	✓	✓	=	=				
	49	✓	✓	✓	=	=				
	50	✓	✓	✓	=	=				
	51	✓	✓	✓	=	=				
	52	✓	✓	✓	=	=				
	53	✓	✓	✓	=	=				
	54	✓	✓	✓	=	=				
	55	✓	✓	✓	=	=				
	56	✓	✓	✓	✓	-				
	57	✓	✓	✓	✓	-				
	58	✓	✓	✓	✓	-				
	59	✓	✓	✓	✓	-				
	60	✓	✓	✓	✓	-				
	61	✓	-	-	-	-				
	62	✓	✓	✓	✓	-				
	63	✓	✓	✓	✓	-				
	64	✓	✓	✓	✓	-				
	65	✓	✓	✓	✓	-				
	66	✓	✓	✓	✓	-				
	67	✓	-	-	-	-				
	68	✓	✓	✓	-	-				
	69	✓	✓	✓	-	-				
	70	✓	✓	✓	-	-				
	71	✓	✓	✓	-	-				
	72	✓	✓	✓	-	-				
	73	✓	✓	✓	-	-				
	74	✓	✓	✓	-	-				
	75	✓	✓	✓	-	-				
	76	✓	✓	✓	-	-				
	77	✓	✓	✓	-	-				
	78	✓	✓	✓	-	-				
	79	✓	✓	✓	-	-				
	80	✓	✓	✓	-	-				
	81	✓	✓	✓	-	-				
	82	✓	✓	✓	-	-				
	83	✓	✓	✓	-	-				
	84	✓	✓	✓	-	-				
	85	✓	✓	✓	-	-				
	86	✓	✓	✓	-	-				
	87	✓	✓	✓	-	-				
	88	✓	✓	✓	-	-				
	89	✓	✓	✓	-	-				
	90	✓	✓	✓	✓	-				
	91	✓	✓	✓	✓	-				
	92	✓	✓	✓	✓	-				
	93	✓	✓	✓	✓	-				
	94	✓	✓	✓	✓	-				

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020	09/22/2020				
	95	✓	✓	✓	=	=				
	96	✓	✓	✓	=	=				
	97	✓	✓	✓	=	=				
	98	✓	✓	✓	=	=				
	99	✓	✓	✓	=	=				
	100	✓	✓	✓	-	-				
	101	✓	✓	✓	-	-				
	102	✓	✓	✓	-	-				
	103	✓	✓	✓	-	-				
	104	✓	✓	✓	-	-				
	105	✓	✓	✓	-	-				
	106	✓	✓	✓	-	-				
	107	✓	✓	✓	-	-				
	108	✓	✓	✓	✓	-				
	109	✓	✓	✓	✓	-				
	110	✓	✓	✓	✓	-				
	111	✓	✓	✓	✓	-				
	112	✓	✓	✓	=	=				
	113	✓	✓	✓	=	=				
	114	✓	✓	✓	=	=				
	115	✓	✓	✓	-	-				
	116	✓	✓	✓	=	=				
	117	✓	✓	✓	✓	=				
	118	✓	✓	✓	✓	=				
	119	✓	✓	✓	✓	=				
	120	✓	✓	✓	✓	=				
	121	✓	✓	✓	✓	=				
	122	✓	✓	✓	✓	=				
	123	✓	✓	✓	✓	=				
	124	✓	✓	✓	✓	=				
	125	✓	✓	✓	✓	=				
	126	✓	✓	✓	✓	=				
	127	✓	✓	✓	✓	=				
	128	✓	✓	✓	-	-				
	129	✓	✓	✓	-	-				
	130	✓	✓	✓	-	-				
	131	✓	✓	✓	-	-				
	132	✓	✓	✓	-	-				
	133	✓	✓	✓	-	-				
	134	✓	✓	✓	-	-				
	135	✓	✓	✓	-	-				
	136	✓	✓	✓	-	-				
	137		✓	✓	✓	-				
	138		✓	✓	✓	-				
	139		✓	✓	✓	-				
	140		✓	✓	✓	-				
	141		✓	✓	✓	-				
	142			✓	=	=				
	143			✓	=	=				
	144			✓	=	=				
	145			✓	-	-				
	146			✓	-	-				

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CLAIM		DATE								
Final	Original	10/10/2018	05/03/2019	12/10/2019	05/13/2020	09/22/2020				
	147			✓	✓	-				
	148			✓	-	-				

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

CPC - Searched*		
Symbol	Date	Examiner
G06F 9/5027	10/10/2018	/WHW/
G06T 1/20	10/10/2018	/WHW/
G06T 1/60	10/10/2018	/WHW/
G06N 99/005	10/10/2018	/WHW/
G06N 3/063	10/10/2018	/WHW/
G06F 2209/509	10/10/2018	/WHW/

CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

/William H. Wood/ Primary Examiner, Art Unit 3992	

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

Search Notes		
Search Notes	**Date**	**Examiner**
Prosecution History Review	10/10/2018	/WHW/
IDS Review	10/10/2018	/WHW/
Litigation Search	10/10/2018	/WHW/
EAST BRS Keyword Search	10/10/2018	/WHW/
NPL IEEE Search	10/10/2018	/WHW/
IDS Review	05/03/2019	/WHW/
EAST BRS	05/03/2019	/WHW/
EAST BRS	12/10/2019	/WHW/
IDS Review	12/10/2019	/WHW/
EAST BRS update	05/13/2020	/WHW/
NPL update	05/13/2020	/WHW/

Interference Search			
US Class/CPC Symbol	**US Subclass/CPC Group**	**Date**	**Examiner**
G06T	1/20; 1/60	09/22/2020	/WHW/
G06F	9/5027; 2209/509	09/22/2020	/WHW/
G06N	3/063; 99/005; 20/00	09/22/2020	/WHW/

/William H. Wood/ Primary Examiner, Art Unit 3992	

NVIDIA Ex1002 Page 1053 of 1081

United States Patent And Trademark Office	**Application No.**	**Art Unit**
	15/808,201	3992
Reissue Terminal Disclaimer Review Form	**Examiner**	**AIA Status**
	WILLIAM H WOOD	No

Original Patent Number of Patent to be Reissued is: 9,189,828	**The Maintenance fee status is:** ◉ up to date ○ not up to date (Consult with SPRS)

Is there a terminal disclaimer filed and accepted during the prosecution of (i) the current reissue application, (ii) the underlying patent, and/or (iii) reexamination proceeding(s) of the underlying patent?

○ NO
◉ YES (Complete the rest of the form)

The reissue patent is subject to Terminal Disclaimer(s) that was/were:

☑ filed and accepted (DISQ or DISQ.E.FILE) during the prosecution of the current reissue application. (Enter terminal disclaimer(s) filing date(s) below).

1. 17 September 2020 2. _____ 3. _____

The underlying patent of the current reissue application is subject to Terminal DIsclaimer(s) that was/were:

☑ accepted (DISQ or DISQ.E.FILE) and of record in the prosecution of the underlying patent and/or reexamination procceeding(s) of the underlying patent. (Enter application/control no(s) and terminal disclaimer(s) filing date(s) below).

1. 14/147,015 30 June 2015 2. _____ _____ 3. _____ _____

EAST Search History

EAST Search History (Interference)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L2	2,860	G06T1/20	USPAT	OR	ON	2020/09/22 17:41
L3	1,835	G06T1/60	USPAT	OR	ON	2020/09/22 17:41
L4	1,309	G06F9/5027	USPAT	OR	ON	2020/09/22 17:41
L5	689	G06N3/063	USPAT	OR	ON	2020/09/22 17:42
L6	216	G06F2209/509	USPAT	OR	ON	2020/09/22 17:42
L7	2,769	G06N99/005	USPAT	OR	ON	2020/09/22 17:42
L8	4,109	G06N20/00	USPAT	OR	ON	2020/09/22 17:43
L9	107	(L2 or L3 or L4 or L5 or L6 or L7 or L8) and (layer near3 (artificial adj neural adj network))	USPAT	OR	ON	2020/09/22 17:46
L10	262	(L2 or L3 or L4 or L5 or L6 or L7 or L8) and (layer with (artificial adj neural adj network))	USPAT	OR	ON	2020/09/22 17:46
L11	357	(L2 or L3 or L4 or L5 or L6 or L7 or L8) and (layer same (artificial adj neural adj network))	USPAT	OR	ON	2020/09/22 17:46
L13	41	(L9 or L10 or L11) and ((transfer$4 or moving or shifting or copying or send$4 or swap$4) near3 (during or while or parallel))	USPAT	OR	ON	2020/09/22 17:47
L14	188	(L9 or L10 or L11) and (layer with neuron)	USPAT	OR	ON	2020/09/22 17:48
L15	176	(L9 or L10 or L11) and (layer near3 neuron)	USPAT	OR	ON	2020/09/22 17:49
L16	144	(L14 or L15) and ((layer or neural or neuron) with (while or during or parallel))	USPAT	OR	ON	2020/09/22 17:49
L17	63	(L14 or L15) and ((layer or neural or neuron) with (simultaneous$4 or concurrent$4))	USPAT	OR	ON	2020/09/22 17:50
L18	162	(L14 or L15) and ((layer or neural or neuron) same (while or during or parallel))	USPAT	OR	ON	2020/09/22 17:50
L19	93	(L14 or L15) and ((input or output) near3 (during or while or simultaneous$4 or concurrent$4 or parallel))	USPAT	OR	ON	2020/09/22 17:50

L20	33	(L16 or L17 or L18 or L19) and (intermediate adj computations)	USPAT	OR	ON	2020/09/22 17:52
L21	0	(L16 or L17 or L18 or L19) and (initializ$4 near2 (texture or shader))	USPAT	OR	ON	2020/09/22 17:55
L22	11	(L16 or L17 or L18 or L19) and (computational adj element)	USPAT	OR	ON	2020/09/22 17:55
L23	0	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) and accelerator) with intelligence)	USPAT	OR	ON	2020/09/22 17:57
L24	1	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) and accelerator) same intelligence)	USPAT	OR	ON	2020/09/22 17:58
L25	14	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) and accelerator) same (neural or neuron))	USPAT	OR	ON	2020/09/22 17:58
L26	9	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) and accelerator) with (neural or neuron))	USPAT	OR	ON	2020/09/22 17:59
L27	0	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) with memory) with intelligence)	USPAT	OR	ON	2020/09/22 18:03
L28	0	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) with memory) same intelligence)	USPAT	OR	ON	2020/09/22 18:03
L29	0	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) with memory) same ((neural or neuron) near3 layer))	USPAT	OR	ON	2020/09/22 18:04
L30	0	(L16 or L17 or L18 or L19) and (((graphics adj processing adj unit) with memory) with ((neural or neuron) near3 layer))	USPAT	OR	ON	2020/09/22 18:04

9/22/2020 6:08:42 PM
C:\Users\wwood\Documents\EAST\Workspaces\15 808201.wsp

Bibliographic Data

Application No: 15/808,201

Foreign Priority claimed: ○ Yes ◉ No

35 USC 119 (a-d) conditions met: ☐ Yes ☑ No ☐ Met After Allowance

Verified and Acknowledged: /William H. Wood/ Examiner's Signature | Initials

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

FILING or 371(c) DATE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.
11/09/2017	345	3992	NRLA-001US03
RULE			

APPLICANTS

Neurala, Inc., Boston, MA,

INVENTORS

Anatoli Gorchetchnikov, Belmont, MA, UNITED STATES

Heather Marie Ames, Milton, MA, UNITED STATES

Massimiliano Versace, Milton, MA, UNITED STATES

Fabrizio Santini, Jamaica Plain, MA, UNITED STATES

CONTINUING DATA

This application is a REI of 14147015 01/03/2014 PAT 9189828

14147015 is a CON of 11860254 09/24/2007 PAT 8648867

11860254 has PRO of 60826892 09/25/2006

FOREIGN APPLICATIONS

IF REQUIRED, FOREIGN LICENSE GRANTED**

12/25/2017

**** SMALL ENTITY ****

STATE OR COUNTRY

UNITED STATES

ADDRESS

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116
UNITED STATES

FILING FEE RECEIVED

$6,160

Corrected Notice of Allowance - Correction of Issue Fee Amount

Due to a technical problem, the Notice of Allowance and Fee(s) Due mailed in this application specified an incorrect fee amount (the amount effective on October 2, 2020 for Notices of Allowance mailed on or after October 2, 2020). Corrected Notices of Allowance are being mailed to all applicants who are affected. This is necessary because the statute requires the applicant to pay the fee that is specified on the Notice of Allowance.

A Corrected Notice of Allowance is attached which specifies the correct issue fee. The time period for payment of the issue fee set forth in the previous Notice of Allowance and Fee(s) Due is being restarted to run THREE MONTHS from the "DATE MAILED" of the attached Corrected Notice of Allowance. Failure to timely comply will result in ABANDONMENT of this application. THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

For applicants who replied prior to receiving the Corrected Notice of Allowance:

Applicant must still file a completed Part B – Fee(s) Transmittal, Form PTOL-85 Part B, or an equivalent, in reply to the Corrected Notice of Allowance since the Corrected Notice of Allowance sets forth a time period for reply. This applies to those situations in which applicant's reply was not processed before the Corrected Notice of Allowance was mailed (e.g., the reply crossed in the mail with the Corrected Notice of Allowance) and those situations in which applicant's reply was already processed. Applicants should check the appropriate box in section 4b of Form PTOL-85 Part B, such as the box authorizing the Director to "charge the required fee(s), any deficiency, or credit any overpayment" to a deposit account number. If this box is checked, the applicant will be refunded the amount of any overpayment to the designated deposit account. In addition, by checking the appropriate box in section 4b of Form PTOL-85, Part B, the applicant requests the Office to "first reapply any previously paid fee shown above," as stated on the form. Applicant may also include an explanatory cover letter, if desired. If, however, applicant prefers to file an equivalent of Form PTOL-85, Part B, and if applicant paid an increased issue fee amount as specified on the Notice of Allowance, applicant may request a refund of the difference between the incorrect amount specified on the Notice of Allowance and the correct amount specified on the Corrected Notice of Allowance.

For questions regarding the issue fee amount calculation, please contact the Application Assistance Unit at 888-786-0101 or *HelpAAUCBuspto.gov*.

For assistance with PAIR or EFS-Web, please contact the Patent EBC at 1-866-217-9197 or *ebc@uspto.gov*.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

CORRECTED
NOTICE OF ALLOWANCE AND FEE(S) DUE

149890 7590 09/30/2020

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	
ART UNIT	PAPER NUMBER
3992	

DATE MAILED: 09/30/2020

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

TITLE OF INVENTION: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
REISSUE	SMALL	$500	$0.00	$0.00	$500	12/30/2020

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. PROSECUTION ON THE MERITS IS CLOSED. THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED. SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE DOES NOT REFLECT A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE IN THIS APPLICATION. IF AN ISSUE FEE HAS PREVIOUSLY BEEN PAID IN THIS APPLICATION (AS SHOWN ABOVE), THE RETURN OF PART B OF THIS FORM WILL BE CONSIDERED A REQUEST TO REAPPLY THE PREVIOUSLY PAID ISSUE FEE TOWARD THE ISSUE FEE NOW DUE.

HOW TO REPLY TO THIS NOTICE:

I. Review the ENTITY STATUS shown above. If the ENTITY STATUS is shown as SMALL or MICRO, verify whether entitlement to that entity status still applies.

If the ENTITY STATUS is the same as shown above, pay the TOTAL FEE(S) DUE shown above.

If the ENTITY STATUS is changed from that shown above, on PART B - FEE(S) TRANSMITTAL, complete section number 5 titled "Change in Entity Status (from status indicated above)".

For purposes of this notice, small entity fees are 1/2 the amount of undiscounted fees, and micro entity fees are 1/2 the amount of small entity fees.

II. PART B - FEE(S) TRANSMITTAL, or its equivalent, must be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted. If an equivalent of Part B is filed, a request to reapply a previously paid issue fee must be clearly made, and delays in processing may occur due to the difficulty in recognizing the paper as an equivalent of Part B.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Maintenance fees are due in utility patents issuing on applications filed on or after Dec. 12, 1980. It is patentee's responsibility to ensure timely payment of maintenance fees when due. More information is available at www.uspto.gov/PatentMaintenanceFees.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), by mail or fax, or via EFS-Web.

By mail, send to: **Mail Stop ISSUE FEE**
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

By fax, send to: **(571)-273-2885**

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

149890 7590 09/30/2020
SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being transmitted to the USPTO via EFS-Web or by facsimile to (571) 273-2885, on the date indicated below.

(Typed or printed name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

TITLE OF INVENTION: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
REISSUE	SMALL	$500	$0.00	$0.00	$500	12/30/2020

EXAMINER	ART UNIT	CLASS-SUBCLASS
WOOD, WILLIAM H	3992	345-531000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

❑ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.

❑ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-09 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list

(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,

(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1 _____

2 _____

3 _____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document must have been previously recorded, or filed for recordation, as set forth in 37 CFR 3.11 and 37 CFR 3.81(a). Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE

(B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent) : ❑ Individual ❑ Corporation or other private group entity ❑ Government

4a. Fees submitted: ❑ Issue Fee ❑ Publication Fee (if required) ❑ Advance Order - # of Copies _____

4b. Method of Payment: *(Please first reapply any previously paid fee shown above)*

❑ Electronic Payment via EFS-Web ❑ Enclosed check ❑ Non-electronic payment by credit card (Attach form PTO-2038)

❑ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment to Deposit Account No. _____

5. **Change in Entity Status** (from status indicated above)

❑ Applicant certifying micro entity status. See 37 CFR 1.29 — NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.

❑ Applicant asserting small entity status. See 37 CFR 1.27 — NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.

❑ Applicant changing to regular undiscounted fee status. — NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature _____ Date _____

Typed or printed name _____ Registration No. _____

PTOL-85 Part B (08-18) Approved for use through 01/31/2020 OMB 0651-0033 U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

NVIDIA Ex1002 Page 1061 of 1081

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 09/30/2020

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

DATE MAILED: 09/30/2020

Determination of Patent Term Extension or Adjustment under 35 U.S.C. 154 (b)

A reissue patent is for "the unexpired part of the term of the original patent." See 35 U.S.C. 251. Accordingly, the above-identified reissue application is not eligible for Patent Term Extension or Adjustment under 35 U.S.C. 154(b).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

OMB Clearance and PRA Burden Statement for PTOL-85 Part B

The Paperwork Reduction Act (PRA) of 1995 requires Federal agencies to obtain Office of Management and Budget approval before requesting most types of information from the public. When OMB approves an agency request to collect information from the public, OMB (i) provides a valid OMB Control Number and expiration date for the agency to display on the instrument that will be used to collect the information and (ii) requires the agency to inform the public about the OMB Control Number's legal significance in accordance with 5 CFR 1320.5(b).

The information collected by PTOL-85 Part B is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450. Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

United States Patent and Trademark Office

Office of the Chief Financial Officer

Document Code:WFEE

User :72230

Sale Accounting Date:09/30/2020

Sale Item Reference Number	Effective Date
15808201	09/17/2020

Document Number	Fee Code	Fee Code Description	Amount Paid	Payment Method
I20209TC44228622	2814	STATUTORY OR TERMINAL DISCLAIMER	$160.00	Deposit Account

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

CPC						
Symbol					**Type**	**Version**
G06T	/	1	/	20	F	2013-01-01
G06F	/	9	/	5027	I	2013-01-01
G06T	/	1	/	60	I	2013-01-01
G06N	/	3	/	063	A	2013-01-01
G06N	/	20	/	00	A	2019-01-01
G06F	/	2209	/	509	A	2013-01-01

CPC Combination Sets							
Symbol					**Type**	**Set**	**Ranking**	**Version**
	/		/					

NONE (Assistant Examiner) (Date)	**Total Claims Allowed:** 57	
/William H. Wood/ Primary Examiner, Art Unit 3992 (Primary Examiner) (Date)	O.G. Print Claim(s) 1	O.G. Print Figure 4

U.S. Patent and Trademark Office Part of Paper No.: 20201013

NVIDIA Ex1002 Page 1065 of 1081

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	Art Unit
	WILLIAM H WOOD	3992

INTERNATIONAL CLASSIFICATION				
CLAIMED				
G06T	/	1	/	60
	/		/	

NON-CLAIMED				
G06F	/	9	/	50
G06T	/	1	/	20

US ORIGINAL CLASSIFICATION	
CLASS	SUBCLASS
1	1

CROSS REFERENCES(S)						
CLASS	SUBCLASS (ONE SUBCLASS PER BLOCK)					

NONE (Assistant Examiner) (Date)	Total Claims Allowed: 57	
/William H. Wood/ Primary Examiner, Art Unit 3992 (Primary Examiner) (Date)	O.G. Print Claim(s) 1	O.G. Print Figure 4

U.S. Patent and Trademark Office Part of Paper No.: 20201013

NVIDIA Ex1002 Page 1066 of 1081

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/808,201	Gorchetchnikov et al.
	Examiner	**Art Unit**
	WILLIAM H WOOD	3992

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☑ T.D. ☐ R.1.47

CLAIMS

Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original
1	1	20	20		39		58		77	36	96		115		134
2	2		21		40		59		78	37	97	43	116		135
3	3		22		41		60		79	38	98	44	117		136
4	4		23	21	42		61		80	39	99	45	118		137
5	5		24	22	43		62		81		100	46	119		138
6	6		25	23	44		63		82		101	47	120		139
7	7		26	24	45		64		83		102	48	121		140
8	8		27	25	46		65		84		103	49	122		141
9	9		28	26	47		66		85		104	50	123	55	142
10	10		29	27	48		67		86		105	51	124	56	143
11	11		30	28	49		68		87		106	52	125	57	144
12	12		31	29	50		69		88		107	53	126		145
13	13		32	30	51		70		89		108	54	127		146
14	14		33	31	52		71		90		109		128		147
15	15		34	32	53		72		91		110		129		148
16	16		35	33	54		73		92		111		130		
17	17		36	34	55		74		93	40	112		131		
18	18		37		56		75		94	41	113		132		
19	19		38		57		76	35	95	42	114		133		

NONE (Assistant Examiner) (Date)	**Total Claims Allowed:** 57
/William H. Wood/ Primary Examiner, Art Unit 3992 (Primary Examiner) (Date)	O.G. Print Claim(s): 1 O.G. Print Figure: 4

U.S. Patent and Trademark Office

Part of Paper No.: 20201013

NVIDIA Ex1002 Page 1067 of 1081

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), by mail or fax, or via EFS-Web.

By mail, send to: Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

By fax, send to: (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

149890 7590 09/30/2020

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being transmitted to the USPTO via EFS-Web or by facsimile to (571) 273-2885, on the date below.

(Typed or printed name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

TITLE OF INVENTION: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$600	$0.00	$0.00	$600	12/30/2020

EXAMINER	ART UNIT	CLASS-SUBCLASS
WOOD, WILLIAM H	3992	345-531000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

❑ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.

❑ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-09 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list
(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,
(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1 Smith Baluch LLP
2
3

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document must have been previously recorded, or filed for recordation, as set forth in 37 CFR 3.11 and 37 CFR 3.81(a). Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE: Neurala, Inc.

(B) RESIDENCE: (CITY and STATE OR COUNTRY): Boston MA

Please check the appropriate assignee category or categories (will not be printed on the patent) : ❑ Individual ☒ Corporation or other private group entity ❑ Government

4a. Fees submitted: ☒ Issue Fee ❑ Publication Fee (if required) ❑ Advance Order - # of Copies _____

4b. Method of Payment: *(Please first reapply any previously paid fee shown above)*

☒ Electronic Payment via EFS-Web ❑ Enclosed check ❑ Non-electronic payment by credit card (Attach form PTO-2038)

☒ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment to Deposit Account No. 60-1876

5. **Change in Entity Status** (from status indicated above)

❑ Applicant certifying micro entity status. See 37 CFR 1.29

❑ Applicant asserting small entity status. See 37 CFR 1.27

❑ Applicant changing to regular undiscounted fee status.

NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.
NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.
NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature /Christopher Max Colice/ Date December 29, 2020

Typed or printed name Christopher Max Colice Registration No. 65,634

Electronic Patent Application Fee Transmittal

Application Number:	15808201
Filing Date:	09-Nov-2017
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Filer:	Christopher M. Colice
Attorney Docket Number:	NRLA-001US03

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
REISSUE ISSUE FEE	2511	1	500	500

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension-of-Time:				
Miscellaneous:				
			Total in USD ($)	**500**

Electronic Acknowledgement Receipt

EFS ID:	41507243
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice
Filer Authorized By:	
Attorney Docket Number:	NRLA-001US03
Receipt Date:	29-DEC-2020
Filing Date:	09-NOV-2017
Time Stamp:	10:58:45
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$500
RAM confirmation Number	E2020BSA59392181
Deposit Account	601876
Authorized User	Christopher Colice

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

37 CFR 1.16 (National application filing, search, and examination fees)

37 CFR 1.17 (Patent application and reexamination processing fees)

37 CFR 1.19 (Document supply fees)

37 CFR 1.20 (Post Issuance fees)

37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Issue Fee Payment (PTO-85B)	NRLA-001US03_Issue_Fee_Payment.pdf	184044 4bf5e45e1e198735a02981bf0f0aca28ce795faa	no	1
Warnings:					
Information:					
2	Fee Worksheet (SB06)	fee-info.pdf	30204 56afcacde013bbb42159170fd2f77de8493dc852	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	214248		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchetchnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT AFTER ALLOWANCE

Amendments to the Specification begin on page 2 of this paper.

Remarks begin on page 3 of this paper.

IN THE SPECIFICATION:

Please replace the paragraph immediately following the heading "Related Applications" with the following paragraph. The amendments to this paragraph are indicated by underlining and brackets.

The present application is a broadening reissue application of U.S. Pat. No. 9,189,828, filed Jan. 3, 2014, which claims a priority benefit, under 35 U.S.C. §120, as a continuation of U.S. application Ser. No. 11/860,254, now U.S. Pat. No. 8,648,867 B2, filed Sep. 24, 2007, entitled "Graphic Processor Based Accelerator System and Method," which in turn claims the priority benefit, under 35 U.S.C. §119(e), of U.S. Application No. 60/826,892, filed Sep. 25, 2006. Each of the above-identified applications is incorporated herein by reference in its entirety. More than one reissue application has been filed for the reissue of U.S. Pat. No. 9,189,828, including this application and a reissue continuation application filed December 29, 2020.

REMARKS

Per MPEP 1451, Applicant is amending the specification to indicate that this application is a reissue application and to include a reference to the reissue continuation application being filed today and claiming priority to this reissue application. This amendment should not affect the claims or the allowability of this reissue application.

The Director is hereby authorized to charge any appropriate fees under 37 C.F.R. §§1.16, 1.17, and 1.21 that may be required by this paper and to credit any overpayment, to Deposit Account No. 60-1876. If an extension of time is required to make this reply timely, please charge the appropriate extension of time fee under 37 C.F.R. § 1.17 to Deposit Account No. 60-1876 if not already paid via EFS.

Dated: December 29, 2020

Respectfully submitted,

Smith Baluch LLP
376 Boylston St, Suite 401
Boston MA 02116

Phone: (617) 947-7280

By /Christopher Max Colice/
Christopher Max Colice
Registration No. 65634

Electronic Acknowledgement Receipt

EFS ID:	41508812
Application Number:	15808201
International Application Number:	
Confirmation Number:	1077
Title of Invention:	GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD
First Named Inventor/Applicant Name:	Anatoli Gorchetchnikov
Customer Number:	149890
Filer:	Christopher M. Colice
Filer Authorized By:	
Attorney Docket Number:	NRLA-001US03
Receipt Date:	29-DEC-2020
Filing Date:	09-NOV-2017
Time Stamp:	12:28:02
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		NRLA-001US03_Amendment_After_Allowance_20201229.pdf	178723 ebd8ac73041b985f09116ee530a949631476c724	yes	3

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Amendment after Notice of Allowance (Rule 312)	1	1
	Specification	2	2
	Applicant Arguments/Remarks Made in an Amendment	3	3
Warnings:			
Information:			
	Total Files Size (in bytes):	178723	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	11/09/2017	Anatoli Gorchetchnikov	NRLA-001US03	1077

149890 7590 01/11/2021

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

EXAMINER	
WOOD, WILLIAM H	

ART UNIT	PAPER NUMBER
3992	

NOTIFICATION DATE	DELIVERY MODE
01/11/2021	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docketing@smithbaluch.com
eofficeaction@appcoll.com
smith@smithbaluch.com

Response to Rule 312 Communication	Application No.	Applicant(s)	
	15/808,201	Gorchetchnikov et al.	
	Examiner	Art Unit	AIA (FITF) Status
	WILLIAM H WOOD	3992	No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

1. ☑ The amendment filed on 29 December 2020 under 37 CFR 1.312 has been considered, and has been:
   a) ☐ entered.
   b) ☑ entered as directed to matters of form not affecting the scope of the invention.
   c) ☐ disapproved because the amendment was filed after the payment of the issue fee.
      Any amendment filed after the date the issue fee is paid must be accompanied by a petition under 37 CFR 1.313(c)(1) and the required fee to withdraw the application from issue.
   d) ☐ disapproved. See explanation below.
   e) ☐ entered in part. See explanation below.

/William H. Wood/ Primary Examiner, Art Unit 3992	/A.J.K/ Supervisory Patent Examiner, Art Unit 3992

OK TO ENTER: /WHW/

01/05/2021

Docket No. NRLA-001US03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Gorchetchnikov et al.

Application No. 15/808,201

Filed: 2017-11-09

Confirmation No. 1077

Art Unit: 3992

Examiner: William H WOOD

Title: GRAPHIC PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD

AMENDMENT AFTER ALLOWANCE

Amendments to the Specification begin on page 2 of this paper.

Remarks begin on page 3 of this paper.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	ISSUE DATE	PATENT NO.	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/808,201	02/16/2021	RE48438	NRLA-001US03	1077

149890 7590 01/27/2021

SMITH BALUCH LLP
376 BOYLSTON ST
STE 401
BOSTON, MA 02116

ISSUE NOTIFICATION

The projected patent number and issue date are specified above.

Determination of Patent Term Extension or Adjustment under 35 U.S.C. 154 (b)

A reissue patent is for "the unexpired part of the term of the original patent." See 35 U.S.C. 251. Accordingly, the above-identified reissue application is not eligible for Patent Term Extension or Adjustment under 35 U.S.C. 154(b).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Application Assistance Unit (AAU) of the Office of Data Management (ODM) at (571)-272-4200.

APPLICANT(s) (Please see PAIR WEB site http://pair.uspto.gov for additional applicants):

Anatoli Gorchetchnikov, Belmont, MA;
Neurala, Inc., Boston, MA, Assignee (with 37 CFR 1.172 Interest);
Heather Marie Ames, Milton, MA;
Massimiliano Versace, Milton, MA;
Fabrizio Santini, Jamaica Plain, MA;

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The USA offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to encourage and facilitate business investment. To learn more about why the USA is the best country in the world to develop technology, manufacture products, and grow your business, visit SelectUSA.gov.