



(12) **United States Patent**
McNamara et al.

(10) **Patent No.:** **US 11,283,261 B2**
(45) **Date of Patent:** ***Mar. 22, 2022**

(54) **MANAGING QUEUE DISTRIBUTION
BETWEEN CRITICAL DATACENTER AND
FLEXIBLE DATACENTER**

(71) Applicant: **Lancium LLC**, Houston, TX (US)

(72) Inventors: **Michael T. McNamara**, Newport
Beach, CA (US); **David J. Henson**,
Houston, TX (US); **Raymond E. Cline,
Jr.**, Houston, TX (US)

(73) Assignee: **LANCIUM LLC**, Houston, TX (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 95 days.

This patent is subject to a terminal dis-
claimer.

(21) Appl. No.: **16/525,142**

(22) Filed: **Jul. 29, 2019**

(65) **Prior Publication Data**

US 2020/0136387 A1 Apr. 30, 2020

Related U.S. Application Data

(63) Continuation of application No. 16/175,335, filed on
Oct. 30, 2018, now Pat. No. 10,367,353.

(51) **Int. Cl.**
H02J 3/12 (2006.01)
G06F 9/48 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC **H02J 3/12** (2013.01); **G06F 1/3206**
(2013.01); **G06F 1/329** (2013.01); **G06F**
9/4893 (2013.01); **G06F 9/5088** (2013.01);
G06F 9/5094 (2013.01)

(58) **Field of Classification Search**

None

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,142,672 A 8/1992 Johnson et al.
5,367,669 A 11/1994 Holland et al.
(Continued)

FOREIGN PATENT DOCUMENTS

CN 101803148 A 8/2010
CN 103163904 A 6/2013
(Continued)

OTHER PUBLICATIONS

Rahimi, Farrokh, "Using a Transactive Energy Framework", IEEE
Electrification Magazine, Dec. 2016, pp. 23-29.

(Continued)


Primary Examiner — Ranodhi Serrao

(74) *Attorney, Agent, or Firm* — Barnes & Thornburg
LLP

(57) **ABSTRACT**

Systems include one or more critical datacenter connected to
behind-the-meter flexible datacenters. The critical datacen-
ter is powered by grid power and not necessarily collocated
with the flexible datacenters, which are powered "behind the
meter." When a computational operation to be performed at
the critical datacenter is identified and determined that it can
be performed more efficiently or advantageously at a flexible
datacenter, the computational operation is instead obtained
by the flexible datacenters for performance. The critical
datacenter and flexible datacenters preferably share a dedi-
cated communication pathway to enable high-bandwidth,
low-latency, secure data transmissions. A queue system may
be used to organize computational operations waiting for
distribution to either the critical datacenter or the flexible
datacenter.

16 Claims, 14 Drawing Sheets



(51)	Int. Cl.			2003/0023885	A1 *	1/2003	Potter	G06F 1/3209
	G06F 1/3206	(2019.01)						713/300
	G06F 9/50	(2006.01)		2003/0037150	A1	2/2003	Nakagawa	
	G06F 1/329	(2019.01)		2003/0074464	A1	4/2003	Bohrer et al.	
				2004/0117330	A1	6/2004	Ehlers et al.	
				2005/0034128	A1	2/2005	Nagashima et al.	
				2005/0203761	A1	9/2005	Barr et al.	
				2006/0161765	A1	7/2006	Cromer et al.	
(56)	References Cited			2006/0253675	A1	11/2006	Johannes Bloks	
	U.S. PATENT DOCUMENTS			2008/0030078	A1	2/2008	Whitted et al.	
				2008/0094797	A1	4/2008	Coglitore et al.	
	5,913,046	A	6/1999	Barth et al.				
	6,288,456	B1	9/2001	Cratty				
	6,633,823	B2 *	10/2003	Bartone	H02J 3/14			
					702/57			
	6,748,932	B1	6/2004	Sorter et al.				
	7,143,300	B2 *	11/2006	Potter	G06F 1/3203			
					713/323			
	7,278,273	B1	10/2007	Whitted et al.				
	7,376,851	B2	5/2008	Kim				
	7,647,516	B2	1/2010	Ranganathan et al.				
	7,702,931	B2	4/2010	Goodrum et al.				
	7,779,276	B2	8/2010	Bolan et al.				
	7,861,102	B1	12/2010	Ranganathan et al.				
	7,921,315	B2	4/2011	Langgood et al.				
	7,970,561	B2	6/2011	Pfeiffer				
	8,001,403	B2	8/2011	Hamilton et al.				
	8,006,108	B2	8/2011	Brey et al.				
	8,214,843	B2	7/2012	Boss et al.				
	8,260,913	B2	9/2012	Knapp et al.				
	8,374,928	B2	2/2013	Gopisetty et al.				
	8,447,993	B2	5/2013	Greene et al.				
	8,571,820	B2	10/2013	Pfeiffer				
	8,627,123	B2	1/2014	Jain et al.				
	8,639,392	B2	1/2014	Chassin				
	8,700,929	B1	4/2014	Weber et al.				
	8,706,915	B2	4/2014	Duchesneau				
	8,719,223	B2	5/2014	Knapp et al.				
	8,789,061	B2	7/2014	Pavel et al.				
	8,799,690	B2	8/2014	Dawson et al.				
	9,003,211	B2	4/2015	Pfeiffer				
	9,003,216	B2	4/2015	Sankar et al.				
	9,026,814	B2	5/2015	Aasheim et al.				
	9,027,024	B2 *	5/2015	Mick	H04L 41/5054			
					718/1			
	9,143,392	B2 *	9/2015	Duchesneau	H04L 41/0816			
	9,207,993	B2	12/2015	Jain				
	9,218,035	B2	12/2015	Li et al.				
	9,252,598	B2	2/2016	Belady et al.				
	9,282,022	B2	3/2016	Matthews et al.				
	9,416,904	B2	8/2016	Belady et al.				
	9,542,231	B2	1/2017	Khan et al.				
	9,552,234	B2	1/2017	Boldyrev et al.				
	9,618,991	B1	4/2017	Clidaras et al.				
	9,645,596	B1	5/2017	Lee et al.				
	9,994,118	B2	6/2018	Williams et al.				
	10,033,210	B2	7/2018	Peterson et al.				
	10,340,696	B2	7/2019	Paine et al.				
	10,367,353	B1	7/2019	McNamara et al.				
	10,367,535	B2	7/2019	Corse et al.				
	10,444,818	B1	10/2019	McNamara et al.				
	10,452,127	B1	10/2019	McNamara et al.				
	10,452,532	B2	10/2019	McVay et al.				
	10,497,072	B2	12/2019	Hooshmand et al.				
	10,608,433	B1	3/2020	McNamara et al.				
	10,618,427	B1	4/2020	McNamara et al.				
	10,637,353	B2	4/2020	Ohyama et al.				
	10,795,428	B2	10/2020	Walsh				
	10,822,992	B2	11/2020	Spears				
	10,862,307	B2	12/2020	Cavness et al.				
	10,873,211	B2	12/2020	McNamara et al.				
	10,931,117	B2	2/2021	Shoemaker				
	11,016,458	B2	5/2021	McNamara et al.				
	11,016,553	B2	5/2021	McNamara et al.				
	11,025,060	B2	6/2021	McNamara et al.				
	11,031,787	B2	6/2021	McNamara et al.				
	2002/0072868	A1 *	6/2002	Bartone	H02J 3/14			
					702/62			
	2002/0158749	A1	10/2002	Ikeda et al.				
				2003/0073150	A1	2/2003	Nakagawa	
				2003/0074464	A1	4/2003	Bohrer et al.	
				2004/0117330	A1	6/2004	Ehlers et al.	
				2005/0034128	A1	2/2005	Nagashima et al.	
				2005/0203761	A1	9/2005	Barr et al.	
				2006/0161765	A1	7/2006	Cromer et al.	
				2006/0253675	A1	11/2006	Johannes Bloks	
				2008/0030078	A1	2/2008	Whitted et al.	
				2008/0094797	A1	4/2008	Coglitore et al.	
				2009/0055665	A1	2/2009	Maglione et al.	
				2009/0070611	A1	3/2009	Bower, III et al.	
				2009/0078401	A1	3/2009	Cichanowicz	
				2009/0089595	A1	4/2009	Brey et al.	
				2009/0216910	A1	8/2009	Duchesneau	
				2010/0211810	A1	8/2010	Zacho	
				2010/0235004	A1	9/2010	Thind	
				2010/0280675	A1	11/2010	Tate, Jr. et al.	
				2010/0328849	A1	12/2010	Ewing et al.	
				2010/0333113	A1	12/2010	Johnson et al.	
				2011/0072289	A1	3/2011	Kato	
				2011/0238342	A1	9/2011	Pfeiffer	
				2012/0000121	A1	1/2012	Swann	
				2012/0072745	A1	3/2012	Ahluwalia et al.	
				2012/0300524	A1	11/2012	Fornage et al.	
				2012/0306271	A1	12/2012	Kuriyama	
				2012/0324259	A1 *	12/2012	Aasheim	G06F 1/329
								713/320
				2012/0326511	A1	12/2012	Johnson	
				2013/0006401	A1	1/2013	Shan	
				2013/0054987	A1	2/2013	Pfeiffer et al.	
				2013/0063991	A1	3/2013	Xiao et al.	
				2013/0086404	A1	4/2013	Sankar et al.	
				2013/0117621	A1 *	5/2013	Saraiya	H04L 49/70
								714/748
				2013/0187464	A1	7/2013	Smith et al.	
				2013/0227139	A1	8/2013	Suffling	
				2013/0304903	A1 *	11/2013	Mick	H04L 67/10
								709/224
				2013/0306276	A1 *	11/2013	Duchesneau	F28D 15/0266
								165/104.21
				2013/0328395	A1	12/2013	Krizman et al.	
				2014/0070756	A1	3/2014	Kearns et al.	
				2014/0137468	A1	5/2014	Ching	
				2014/0180886	A1	6/2014	Forbes, Jr.	
				2014/0379156	A1	12/2014	Kamel et al.	
				2015/0012113	A1	1/2015	Celebi	
				2015/0121113	A1	4/2015	Ramamurthy et al.	
				2015/0155712	A1	6/2015	Mondal	
				2015/0212122	A1	7/2015	Sobotka et al.	
				2015/0229227	A1	8/2015	Aeloiza et al.	
				2015/0277410	A1	10/2015	Gupta et al.	
				2015/0278968	A1 *	10/2015	Steven	G06Q 10/06314
								705/7.35
				2015/0288183	A1	10/2015	Villanueva et al.	
				2015/0372538	A1	12/2015	Siegler et al.	
				2016/0006066	A1	1/2016	Robertson	
				2016/0011617	A1	1/2016	Liu et al.	
				2016/0043552	A1	2/2016	Villanueva, Jr. et al.	
				2016/0087909	A1	3/2016	Chatterjee et al.	
				2016/0126783	A1	5/2016	Cheng et al.	
				2016/0170469	A1	6/2016	Sehgal et al.	
				2016/0172900	A1	6/2016	Welch, Jr.	
				2016/0187906	A1	6/2016	Bodas et al.	
				2016/0198656	A1	7/2016	McNamara et al.	
				2016/0202744	A1	7/2016	Castro-Leon	
				2016/0212954	A1	7/2016	Argento	
				2016/0248631	A1	8/2016	Duchesneau	
				2016/0324077	A1	11/2016	Frantzen et al.	
				2017/0023969	A1	1/2017	Shows et al.	
				2017/0104336	A1	4/2017	Eibsat et al.	
				2017/0261949	A1	9/2017	Hoffmann et al.	
				2017/0373500	A1	12/2017	Shafi et al.	
				2018/0026478	A1	1/2018	Peloso	
				2018/0116070	A1	4/2018	Broadbent et al.	
				2018/0144414	A1	5/2018	Lee et al.	
				2018/0202825	A1	7/2018	You et al.	
				2018/0240112	A1	8/2018	Castinado et al.	

(56)

References Cited**U.S. PATENT DOCUMENTS**

2018/0366978	A1	12/2018	Matan et al.
2018/0367320	A1	12/2018	Montalvo
2019/0052094	A1	2/2019	Pmsvsv et al.
2019/0168630	A1	6/2019	Mrlik et al.
2019/0258307	A1	8/2019	Shaikh et al.
2019/0280521	A1	9/2019	Lundstrom et al.
2019/0318327	A1	10/2019	Sowell et al.
2019/0324820	A1	10/2019	Krishnan et al.
2020/0040272	A1	2/2020	Cavness et al.
2020/0051184	A1	2/2020	Barbour
2020/0073466	A1	3/2020	Walsh
2020/0136388	A1	4/2020	McNamara et al.
2020/0167197	A1	5/2020	Bahramshahry et al.
2021/0175710	A1	6/2021	Campbell et al.
2021/0287309	A1	9/2021	Gebhardt et al.
2021/0298195	A1	9/2021	Barbour

FOREIGN PATENT DOCUMENTS

CN	103748757	A	4/2014
CN	104969434	A	10/2015
KR	20090012523	A	2/2009
WO	2015/199629		12/2015
WO	WO-2019116375	A1	6/2019
WO	WO-2019139633	A1	7/2019
WO	WO-2020227811	A1	11/2020

OTHER PUBLICATIONS

Ghamkhari et al., "Optimal Integration of Renewable Energy Resources in Data Centers with Behind-the-Meter Renewable Generator," Department of Electrical and Computer Engineering Texas Tech University, 2012, pp. 3340-3444.

Non-Final Office Action dated Dec. 5, 2019 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 72 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 72 pages.

Non-Final Office Action dated Nov. 14, 2019 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 25 pages.

Non-Final Office Action dated Nov. 21, 2019 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 57 pages.

Non-Final Office Action dated Dec. 11, 2019 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 17 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/528,348, filed Oct. 8, 2019, 33 pages.

Notice of Allowance dated Aug. 15, 2019, for U.S. Appl. No. 16/175,146, filed Oct. 30, 2018, 17 pages.

Co-Pending U.S. Appl. No. 16/132,092, filed Sep. 14, 2018.

Co-Pending U.S. Appl. No. 16/175,246, filed Oct. 30, 2018.

Co-Pending U.S. Appl. No. 16/482,495, filed Jul. 31, 2019.

Notice of Allowance dated Jan. 27, 2020, for U.S. Appl. No. 16/702,931, filed Dec. 4, 2019, 23 pages.

Notice of Allowance dated Apr. 2, 2019, for U.S. Appl. No. 16/175,335, filed Oct. 30, 2018, 12 pages.

Notice of Allowance dated Jul. 29, 2019, for U.S. Appl. No. 16/245,532, filed Jan. 11, 2019, 13 pages.

Advisory Action dated Nov. 13, 2020 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 182 pages.

Advisory Action dated Oct. 22, 2020 for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 3 pages.

Bakar.N et al., "Microgrid and Load Shedding Scheme During Islanded Mode: a Review," Elsevier, May 26, 2020, vol. 71, pp. 161-169. <https://www.sciencedirect.com/science/article/pii/S1364032116311030>.

Choi.Y et al., "Optimal Load Shedding for Maximizing Satisfaction in an Islanded Microgrid," Energies, 2017, vol. 10, p. 45, doi: 10.3390/en10010045.

Final Office Action dated Jul. 23, 2020 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 26 pages.

Final Office Action dated May 19, 2020 for U.S. Appl. No. 16/809,111, filed Mar. 4, 2020, 36 pages.

Final Office Action dated Jun. 3, 2020 for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 33 pages.

Final Office Action dated May 28, 2020 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 24 pages.

Final Office Action dated Jul. 29, 2020 for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 5 pages.

Gao.H et al., "Dynamic Load Shedding for an Islanded Microgrid With Limited Generation Resources," IET Generation, Transmission & Distribution, Sep. 2016, vol. 10(12), pp. 2953-2961. doi: 10.1049/iet-gtd.2015.1452.

International Search Report and Written Opinion of PCT Application No. PCT/US2020/044536, dated Aug. 26, 2020, 24 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US2020/044539, dated Aug. 26, 2020, 7 pages.

John, "Stem and CPower to Combine Behind-the-Meter Batteries and Demand Response," Energy Storage, Aug. 8, 2017, 1 page.

Lim.Y et al., "Distributed Load-shedding System for Agent-based Autonomous Microgrid Operations," Energies, 2014, vol. 7(1), pp. 385-401. doi: 10.3390/en7010385.

Liu.W et al., "Improved Average Consensus Algorithm Based Distributed Cost Optimization for Loading Shedding of Autonomous Microgrids," International Journal of Electrical Power & Energy Systems, Dec. 2015, vol. 73, pp. 89-96. doi: 10.1016/j.ijepes.2015.04.006.

Mousavizadeh.S et al., "A Linear Two-stage Method for Resiliency Analysis in Distribution Systems Considering Renewable Energy and Demand Response Resources," Elsevier, 2017, pp. 443-460. doi: 10.1016/j.apenergy.2017.11.067.

Non-Final Office Action dated Jun. 12, 2020 for U.S. Appl. No. 16/803,109, filed Dec. 27, 2020, 31 pages.

Non-Final Office Action dated May 14, 2020 on for U.S. Appl. No. 16/834,987, filed Mar. 30, 2020, 30 pages.

Notice of Allowance dated Oct. 13, 2020 on for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 5 pages.

Notice of Allowance dated Jun. 12, 2020 on for U.S. Appl. No. 16/834,987, filed Mar. 30, 2020, 9 pages.

Notice of Allowance dated Sep. 17, 2020 on for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 5 pages.

Notice of Allowance dated Nov. 19, 2020 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 7 pages.

Notice of Allowance dated Jul. 29, 2020 on for U.S. Appl. No. 16/132,011, filed Sep. 14, 2018, 5 pages.

Notice of Allowance dated Oct. 29, 2020 on for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 8 pages.

Pashajavid.E et al., "A Multimode Supervisory Control Scheme for Coupling Remote Droop-Regulated Microgrids," IEEE Transactions on Smart Grid, May 26, 2020, vol. 9(5), pp. 5381-5392. <https://ieeexplore.ieee.org/abstract/document/7888570/>.

Pashajavid.E et al., "Frequency Support for Remote Microgrid Systems With Intermittent Distributed Energy Resources—A Two-level Hierarchical Strategy," IEEE Systems Journal, May 26, 2020, vol. 12(3), pp. 2760-2771. <https://ieeexplore.ieee.org/abstract/document/7862156/>.

Rudez.U and Mihalic.R, "Predictive Underfrequency Load Shedding Scheme for Islanded Power Systems With Renewable Generation," Electric Power Systems Research, May 2015, vol. 126, pp. 21-28. doi: 10.1016/j.epsr.2015.04.017.

Wilson, Joseph Nathanael, "A Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University, 2016, 154 pages.

Xu.Q et al., "Distributed Load Shedding for Microgrid With Compensation Support via Wireless Network," IET Generation, Transmission & Distribution, May 2018, vol. 12(9), pp. 2006-2018. doi: 10.1049/iet-gtd.2017.1029.

Zhou.Q et al., "Two-Stage Load Shedding for Secondary Control in Hierarchical Operation of Islanded Microgrids," IEEE Transactions on Smart Grid, May 2019, vol. 10(3), pp. 3103-3111. doi: 10.1109/TSG.2018.2817738.

(56)

References Cited**OTHER PUBLICATIONS**

Final Office Action dated Oct. 1, 2019 for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 18 pages.

Wilson, J., "A Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University, 2016.

Final Office Action dated Apr. 17, 2020 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 59 pages.

Hayes, Adam S., "A Cost of Production Model for Bitcoin," Department of Economics, The New School for Social Research, Mar. 2015, 5 pages.

Non-Final Office Action dated Feb. 20, 2020 for U.S. Appl. No. 16/702,894, filed Dec. 4, 2019, 30 pages.

Non-Final Office Action dated Apr. 2, 2020 on for U.S. Appl. No. 16/132,011, filed Sep. 14, 2018, 5 pages.

Non-Final Office Action dated Mar. 30, 2020 for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 46 pages.

Notice of Allowance dated Mar. 2, 2020, for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 15 pages.

Notice of Allowance dated Apr. 6, 2020, for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 12 pages.

Bird et al., "Wind and Solar Energy Curtailment: Experience and Practices in the United States," NREL, Technical Report NREL/TP-6A20-60983, Mar. 2014, 58 pages.

EPEX Spot, "How They Occur, What They Mean," https://www.epeexspot.com/en/company-info/basics_of_the_power_market/negative_prices, 2018, 2 pages.

International Searching Authority, International Search Report and Written Opinion, dated May 31, 2018, issued in connection with International Patent Application No. PCT/US2018/017950, filed on Feb. 13, 2018, 15 pages.

International Searching Authority, International Search Report and Written Opinion, dated Apr. 30, 2018, issued in connection with International Patent Application No. PCT/US2018/017955, filed on Feb. 13, 2018, 22 pages.

Soluna, Powering the Blockchain, Aug. 2018, Verison 1.1, 29 pages.

Non-Final Office Action dated Feb. 4, 2021 on for U.S. Appl. No. 16/284,610, filed Feb. 25, 2019, 9 pages.

Non-Final Office Action dated Feb. 12, 2021 on for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 25 pages.

Notice of Allowance dated Feb. 8, 2021 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 21 pages.

Notice of Allowance dated Feb. 8, 2021 on for U.S. Appl. No. 16/803,109, filed Feb. 27, 2020, 29 pages.

Notice of Allowance dated Feb. 8, 2021 on for U.S. Appl. No. 16/834,987, filed Mar. 30, 2020, 180 pages.

Notice of Allowance dated Jan. 13, 2021 on for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 5 pages.

Notice of Allowance dated Jan. 25, 2021 on for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 5 pages.

Notice of Allowance dated Jan. 25, 2021 on for U.S. Appl. No. 16/702,894, filed Dec. 4, 2019, 24 pages.

Notice of Allowance dated Jan. 27, 2021 on for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 8 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US2021/019875, dated Apr. 29, 2021, 12 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US20/57686, dated Feb. 22, 2021, 67 pages.

Iso, "Distributed Energy Resources Roadmap for New York's Wholesale Electricity Markets," Distributed Energy Resource Roadmap, Jan. 2017, pp. 1-39. [retrieved on Dec. 15, 2020], Retrieved from the Internet: <url:="https://www.nyiso.com/documents/20142/1391862/Distributed_Energy_Resources_Roadmap.pdf/ec0b3b64-4de2-73e0-ffef-49a4b8b1">https://www.nyiso.com/documents/20142/1391862/Distributed_Energy_Resources_Roadmap.pdf/ec0b3b64-4de2-73e0-ffef-49a4b8b1 b3ca.</url:>.

Kewl, "Start-Up From the Heart of Berlin Has Pioneered Decentralized Mobile Mining by Combining Blockchain With Regenerative Energy" Nov. 2017, 3 pages pdf <http://www.crypto-news.net/start-up-from-the-heart-of-berlin-has-pioneered-decentralized-mobile-mining-by-combining-blockchain-with-regenerative-energy/> (Year: 2017).

Non-Final Office Action dated Apr. 1, 2021 on for U.S. Appl. No. 16/482,495, filed Jul. 31, 2019, 59 pages.

Non-Final Office Action dated Mar. 25, 2021 on for U.S. Appl. No. 16/573,577, filed Sep. 17, 2019, 65 pages.

Non-Final Office Action dated May 11, 2021 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 64 pages.

Non-Final Office Action dated May 28, 2021 for U.S. Appl. No. 16/658,983, filed Oct. 21, 2019, 21 pages.

Notice of Allowance dated May 12, 2021 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 2 pages.

Notice of Allowance dated Apr. 20, 2021 on for U.S. Appl. No. 16/482,495, filed Jul. 31, 2019, 5 pages.

Notice of Allowance dated May 27, 2021 on for U.S. Appl. No. 16/284,610, filed Feb. 25, 2019, 16 pages.

Notice of Allowance dated May 27, 2021 on for U.S. Appl. No. 16/284,610, filed Feb. 25, 2019, 5 pages.

European Patent Application No. EP18900411.2, Partial Supplementary European Search Report dated Sep. 9, 2021.

Final Office Action dated Aug. 9, 2021 on for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 43 pages.

Final Office Action dated Aug. 9, 2021 on for U.S. Appl. No. 16/573,577, filed Sep. 17, 2019, 16 pages.

Ghatikar et al., "Demand Response Opportunities and Enabling Technologies for DataCenters: Findings from Field Studies," Lawrence Berkeley National Laboratory, Aug. 2012, 57 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US2021/045972, dated Nov. 15, 2021, 16 pages.


Notice of Allowance dated Jun. 9, 2021 for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 11 pages.


Notice of Allowance dated Jul. 26, 2021 on for U.S. Appl. No. 16/284,610, filed Feb. 25, 2019, 2 pages.

Notice of Allowance dated Oct. 8, 2021 on for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 3 pages.

Notice of Allowance dated Sep. 24, 2021 for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 06 pages.

* cited by examiner

**FIG. 1**








FIG. 3

**FIG. 4**

**FIG. 5**




FIG. 6







FIG. 7

**FIG. 8**

**FIG. 9A**

**FIG. 9B**







FIG. 10

**FIG. 11**

**FIG. 12**

**FIG. 13**

1

MANAGING QUEUE DISTRIBUTION BETWEEN CRITICAL DATACENTER AND FLEXIBLE DATACENTER

CROSS-REFERENCE TO RELATED APPLICATION

The present application is a continuation of U.S. patent application Ser. No. 16/175,335, filed on Oct. 30, 2018, the entire contents of which are herein incorporated by reference.

FIELD OF THE INVENTION

This specification relates to a system for controlling the use of “behind-the-meter” power.

BACKGROUND OF THE INVENTION

The price for power distributed through regional and national electric power grids is composed of Generation, Administration, and Transmission & Distribution (“T&D”) costs. T&D costs are a significant portion of the overall price paid by consumers for electricity. T&D costs include capital costs (land, equipment, substations, wire, etc.), electrical transmission losses, and operation and maintenance costs. Electrical power is typically generated at local stations (e.g., coal, natural gas, nuclear, and renewable sources) in the Medium Voltage class of 2.4 kVAC to 69 kVAC before being converted in an AC-AC step up transformer to High Voltage at 115 kVAC or above. T&D costs are accrued at the point the generated power leaves the local station and is converted to High Voltage electricity for transmission onto the grid.

Local station operators are paid a variable market price for the amount of power leaving the local station and entering the grid. However, grid stability requires that a balance exist between the amount of power entering the grid and the amount of power used from the grid. Grid stability and congestion is the responsibility of the grid operator and grid operators take steps, including curtailment, to reduce power supply from local stations when necessary. Frequently, the market price paid for generated power will be decreased in order to disincentivize local stations from generating power. In some cases, the market price will go negative, resulting in a cost to local station operators who continue to supply power onto a grid. Grid operators may sometimes explicitly direct a local station operator to reduce or stop the amount of power the local station is supplying to the grid.

Power market fluctuations, power system conditions such as power factor fluctuation or local station startup and testing, and operational directives resulting in reduced or discontinued generation all can have disparate effects on renewal energy generators and can occur multiple times in a day and last for indeterminate periods of time. Curtailment, in particular, is particularly problematic.

According to the National Renewable Energy Laboratory’s Technical Report TP-6A20-60983 (March 2014):

[C]urtailment [is] a reduction in the output of a generator from what it could otherwise produce given available resources (e.g., wind or sunlight), typically on an involuntary basis. Curtailments can result when operators or utilities command wind and solar generators to reduce output to minimize transmission congestion or otherwise manage the system or achieve the optimal mix of resources. Curtailment of wind and solar resources typically occurs because of transmission con-

2

gestion or lack of transmission access, but it can also occur for reasons such as excess generation during low load periods that could cause baseload generators to reach minimum generation thresholds, because of voltage or interconnection issues, or to maintain frequency requirements, particularly for small, isolated grids. Curtailment is one among many tools to maintain system energy balance, which can also include grid capacity, hydropower and thermal generation, demand response, storage, and institutional changes. Deciding which method to use is primarily a matter of economics and operational practice.

“Curtailment” today does not necessarily mean what it did in the early 2000s. Two sea changes in the electric sector have shaped curtailment practices since that time: the utility-scale deployment of wind power, which has no fuel cost, and the evolution of wholesale power markets. These simultaneous changes have led to new operational challenges but have also expanded the array of market-based tools for addressing them.

Practices vary significantly by region and market design. In places with centrally-organized wholesale power markets and experience with wind power, manual wind energy curtailment processes are increasingly being replaced by transparent offer-based market mechanisms that base dispatch on economics. Market protocols that dispatch generation based on economics can also result in renewable energy plants generating less than what they could potentially produce with available wind or sunlight. This is often referred to by grid operators by other terms, such as “downward dispatch.” In places served primarily by vertically integrated utilities, power purchase agreements (PPAs) between the utility and the wind developer increasingly contain financial provisions for curtailment contingencies.

Some reductions in output are determined by how a wind operator values dispatch versus non-dispatch. Other curtailments of wind are determined by the grid operator in response to potential reliability events. Still other curtailments result from overdevelopment of wind power in transmission-constrained areas.

Dispatch below maximum output (curtailment) can be more of an issue for wind and solar generators than it is for fossil generation units because of differences in their cost structures. The economics of wind and solar generation depend on the ability to generate electricity whenever there is sufficient sunlight or wind to power their facilities.

Because wind and solar generators have substantial capital costs but no fuel costs (i.e., minimal variable costs), maximizing output improves their ability to recover capital costs. In contrast, fossil generators have higher variable costs, such as fuel costs. Avoiding these costs can, depending on the economics of a specific generator, to some degree reduce the financial impact of curtailment, especially if the generator’s capital costs are included in a utility’s rate base.

Curtailment may result in available energy being wasted (which may not be true to the same extent for fossil generation units which can simply reduce the amount of fuel that is being used). With wind generation, in particular, it may also take some time for a wind farm to become fully operational following curtailment. As such, until the time that the wind farm is fully operational, the wind farm may not be operating with optimum efficiency and/or may not be able to provide power to the grid.

3

BRIEF SUMMARY OF THE INVENTION

In an example, a system is described. The system includes a flexible datacenter comprising: a behind-the-meter power input system, a first power distribution system, a datacenter control system, and a first plurality of computing systems powered by the behind-the-meter power input system via the first power distribution system. The flexible datacenter control system is configured to modulate power delivery to the plurality of computing systems based on one or more monitored power system conditions or an operational directive. The system further includes a critical datacenter comprising: a power input system, a second power distribution system, a critical datacenter control system, and a second plurality of computing systems powered by the power input system via the second power distribution system. The system further includes a queue system configured to organize a plurality of computational operations and a first communication link connecting the flexible datacenter, the critical datacenter, and the queue system. The system also includes a routing control system configured to (i) identify, using the queue system, a computational operation to be performed, (ii) determine whether to route the computational operation to the flexible datacenter, and (iii) based on a determination to route the computational operation to the flexible datacenter, cause the computational operation to be sent to the flexible datacenter via the first communication link.

In another example, a system is described. The system includes a plurality of flexible datacenters, each flexible datacenter comprising: a behind-the-meter power input system, a first power distribution system, a datacenter control system, and a first plurality of computing systems powered by the behind-the-meter power input system. The flexible datacenter control system is configured to modulate power delivery to the plurality of computing systems based on one or more monitored power system conditions or an operational directive. The system further includes a critical datacenter comprising: a power input system, a second power distribution system, a critical datacenter control system, and a second plurality of computing systems powered by the power input system via the second power distribution system. The system also includes a queue system configured to organize a plurality of computational operations and a first communication link connecting the plurality of flexible datacenter, the critical datacenter, and the queue system. The system further includes a routing control system configured to (i) identify, using the queue system, a computational operation to be performed, (ii) determine whether to route the computational operation to a flexible datacenter in the plurality of flexible datacenters, (iii) based on a determination to route the computational operation to a flexible datacenter in the plurality of flexible datacenters, determine a specific flexible datacenter in the plurality of flexible datacenters to route the computational operation to, and (iv) cause the computational operation to be sent to the specific flexible datacenter via the first communication link.

Other aspects of the present invention will be apparent from the following description and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a computing system in accordance with one or more embodiments of the present invention.

FIG. 2 shows a flexible datacenter in accordance with one or more embodiments of the present invention.

4

FIG. 3 shows a three-phase power distribution of a flexible datacenter in accordance with one or more embodiments of the present invention.

FIG. 4 shows a control distribution scheme of a flexible datacenter in accordance with one or more embodiments of the present invention.

FIG. 5 shows a control distribution scheme of a fleet of flexible datacenters in accordance with one or more embodiments of the present invention.

FIG. 6 shows a flexible datacenter powered by one or more wind turbines in accordance with one or more embodiments of the present invention.

FIG. 7 shows a flexible datacenter powered by one or more solar panels in accordance with one or more embodiments of the present invention.

FIG. 8 shows a flexible datacenter powered by flare gas in accordance with one or more embodiments of the present invention.

FIG. 9A shows a method of dynamic power delivery to a flexible datacenter using behind-the-meter power in accordance with one or more embodiments of the present invention.

FIG. 9B shows another method of dynamic power delivery to a flexible datacenter using behind-the-meter power in accordance with one or more embodiments of the present invention.

FIG. 10 illustrates a system for managing queue distribution among a critical datacenter and behind-the-meter flexible datacenters in accordance with one or more embodiments of the present invention.

FIG. 11 illustrates a system for managing queue distribution among a critical datacenter and a plurality of behind-the-meter flexible datacenters in accordance with one or more embodiments of the present invention.

FIG. 12 illustrates a method for managing queue distribution between a critical datacenter and a flexible datacenter in accordance with one or more embodiments of the present invention.

FIG. 13 illustrates a method for managing queue distribution between a critical datacenter and a plurality of flexible datacenter in accordance with one or more embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

One or more embodiments of the present invention are described in detail with reference to the accompanying figures. For consistency, like elements in the various figures are denoted by like reference numerals. In the following detailed description of the present invention, specific details are set forth in order to provide a thorough understanding of the present invention. In other instances, well-known features to one having ordinary skill in the art are not described to avoid obscuring the description of the present invention.

The embodiments provided herein relate to providing an electrical load “behind the meter” at local stations such that generated power can be directed to the behind-the-meter load instead of onto the grid, typically for intermittent periods of time. “Behind-the-meter” power includes power that is received from a power generation system (for instance, but not limited to, a wind or solar power generation system) prior to the power undergoing step-up transformation to High Voltage class AC power for transmission to the grid. Behind-the-meter power may therefore include power

5

drawn directly from an intermittent grid-scale power generation system (e.g. a wind farm or a solar array) and not from the grid.

The embodiments herein provide an advantage when, for example, the power system conditions exhibit excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, conditions where the cost for power is economically viable (e.g., low cost for power), or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation. This is not least because the excess power can be utilized by the behind-the-meter electrical load rather than going to waste. In addition, by providing an electrical load behind-the-meter rather than connected to the grid, electrical transmission losses resulting from transmission of power through the grid can be reduced. In addition, any degradation in the power generation systems which may result from curtailment may be reduced.

Preferably, controlled computing systems that consume electrical power through computational operations can provide a behind-the-meter electrical load that can be granularly ramped up and down quickly under the supervision of control systems that monitor power system conditions and direct the power state and/or computational activity of the computing systems. In one embodiment, the computing systems preferably receive all their power for computational operations from a behind-the-meter power source. In another embodiment, the computing systems may additionally include a connection to grid power for supervisory and communication systems or other ancillary needs. In yet another embodiment, the computing systems can be configured to switch between behind-the-meter power and grid power under the direction of a control system.

Among other benefits, a computing system load with controlled granular ramping allows a local station to avoid negative power market pricing and to respond quickly to grid directives. Local stations may include a station capable of controlling power direction and supply and may be referred to as substations or station controls. For instance, a local station may control access to power from the power grid.

Various computing systems can provide granular behind-the-meter ramping. Preferably the computing systems perform computational tasks that are immune to, or not substantially hindered by, frequent interruptions or slow-downs in processing as the computing systems ramp up and down. In one embodiment, control systems can activate or deactivate one or more computing systems in an array of similar or identical computing systems sited behind the meter. For example, one or more blockchain miners, or groups of blockchain miners, in an array may be turned on or off. In another embodiment, control systems can direct time-insensitive computational tasks to computational hardware, such as CPUs and GPUs, sited behind the meter, while other hardware is sited in front of the meter and possibly remote from the behind-the-meter hardware. Any parallel computing processes, such as Monte Carlo simulations, batch processing of financial transactions, graphics rendering, and oil and gas field simulation models are all good candidates for such interruptible computational operations.

A typical datacenter provides computational resources to support computational operations. Particularly, one or more

6

enterprises may assign computational operations to the typical datacenter with expectations that the typical datacenter reliably provides resources to support the computational operations, such as processing abilities, networking, and/or storage. The computational operations assigned to a typical datacenter may vary in their requirements. Some computational operations may require low-latency processing, or are extremely time sensitive, or require a high degree of support and reliability from the datacenter. Other computational operations are not time sensitive and can be batch processed over time, or can be distributed across multiple computational systems with interruptible parallel processing, or can be run on specialized hardware for more efficient processing. Therefore, there can be an economic advantage to sending computational operations to different types of datacenters that have different costs for different types of computational operations. According to embodiments disclosed here, a system of one or more high-compute-cost critical datacenters and one or more low-compute-cost flexible datacenters provides such an economic advantage.

A critical datacenter may have a similar configuration to a typical datacenter. Due to the need to reliably provide computing resources to support critical operations, a critical datacenter is preferably connected to a reliable power source, such as the power grid with multiple redundant power supply systems. The power grid will offer a constant power supply that the critical datacenter uses to meet the needs of assigned computational operations. However, the grid power that enables the critical datacenter to provide the required computational resources is a very significant expense.

In addition, it might also be difficult to estimate future costs associated with utilizing the critical datacenter for critical computational operations. The cost for power from the power grid can fluctuate in price depending on various factors, including the location of the critical datacenter using the power, the overall demand for the power, weather conditions, fuel costs endured by suppliers of the power to the power grid, and time of use, among others.

Example embodiments presented herein aim to reduce the cost associated with using a critical database to perform computational operations. In particular, some examples involve using one or more flexible datacenters to offload computational operations from a critical datacenter. A flexible datacenter may also initially assume and support computational operations rather than a critical datacenter supporting the computational operations. As described below with regards to FIG. 2, a flexible datacenter may use behind-the-meter power in order to provide processing abilities and other computing resources. By using behind-the-meter power from renewable energy sources (e.g., wind farm 600, solar farm 700) and other behind-the-meter power sources, a flexible datacenter can provide computing resources at very low costs, significantly below the costs incurred to power a critical datacenter. As such, one or more flexible datacenters may assist a critical datacenter in efficiently handling computational operations assigned to the critical datacenter by one or more enterprises. The addition of one or more flexible datacenters can also increase the quantity of computational resources available to perform and support computational operations.

In addition, some example implementations presented herein involve the use of queue system. A queue system is a data model that can organize computational operations for subsequent access and distribution to available datacenters, including critical datacenters and flexible datacenters. In various embodiments, one or more control systems may

support and maintain the queue system. For example, the remote master control system **420**, the local station control system **410**, or the datacenter control system **220** may support operations of queue system. By using the queue system, computational operations can be organized and distributed efficiently to datacenters that have the capabilities and availabilities to handle the computational operations.

The structure and operation of the queue system may vary within examples. Particularly, the queue system may include one or more queues that arrange the computational operations according to a variety of factors. Example factors include parameters of each computational operation, deadlines for completing each computational operation, time of submission of each computational operation, computing resources required for each computational operation, and price obtained to perform each computational operation, among others. As such, the queue system may organize and maintain the computational operations waiting for performance by the critical datacenter or the flexible datacenter.

Some examples may involve using a centralized queue system maintained centrally by a control system. For instance, the remote master control system **420** may support and maintain the queue system as a centralized queue. The remote master control system **420** may receive new computational operations requests from enterprises or other sources. The remote master control system **420** may then place each new computational operation request in the centralized queue system for access and distribution to one or more datacenters. In some examples, the remote master control system **420** may manage the distribution of computational operations to available datacenters. The distribution may involve factoring the availability of the datacenters, the type of datacenters (e.g., critical datacenter or flexible datacenter), and the cost for a datacenter to perform the computational operation, among others. In other examples, the data control systems at each datacenter may access and obtain computational operations from the centralized queue system.

Some example embodiments may involve using a decentralized queue system. Particularly, the queue system may be distributed across multiple control systems. For example, a first queue subsystem of the decentralized queue system may receive and organize computational operations for a first set of datacenters to support and a second queue subsystem of the decentralized queue system may receive and organize computational operations for a second set of datacenters to support. The decentralized queue system may require less computational resources to support and maintain. Further, by dividing computational operations into multiple queue subsystems, sets of datacenters may be similarly divided to quickly and efficiently address computational operations within each subsystem.

In various embodiments, the access to computational operations may differ. In some examples, the data control system or computing system managing the queue system may control access to computational operations assigned to the queue system. Particularly, the remote master control system **420** or another control system may communicate and distribute computational operations to critical datacenters and flexible datacenters. In other examples, the control system at each datacenter may access and assume responsibility for computational operations placed within the queue system. This way, the datacenter control system **220** may manage computational operations performed at the flexible datacenter **200** based on capabilities and availability of the computing systems **100** at the flexible datacenter **200**.

In some examples, a critical datacenter may offload some or all of a set of computational operations to the queue system. The critical datacenter may also release some or all of the set of computational operations directly to flexible datacenter. Particularly, when conditions signal that use of a flexible datacenter is economically viable (i.e., at the same or decreased costs relative to using power from the power grid at the critical datacenter), a flexible datacenter may assume some or even all of one or more sets of computational operations from the critical datacenter.

In some instances, a flexible datacenter may assume less critical computational operations from the queue system or directly from a critical datacenter. This way, the critical datacenter may offload less critical computational operations directly or indirectly to a flexible datacenter to support and manage. In such a configuration, the critical datacenter may continue to support critical operations assigned to the critical datacenter by one or more enterprises while offloading less critical operations directly or indirectly to one or more flexible datacenters. As a result, the critical datacenter may ensure that the critical operations remain supported by computational resources powered by grid power.

In other examples, a flexible datacenter may assume critical operations, augmenting the resources provided by the critical datacenter. Particularly, situations can arise where the flexible datacenter can operate at a lower cost than the critical datacenter. For instance, one or more behind-the-meter power sources (e.g., wind farm **600**, solar farm **700**) may enable the flexible datacenter to operate at a lower cost than the critical datacenter. As a result, using the flexible datacenter instead of the critical datacenter can lower the costs required to support assigned computing operations. If the situation changes such that the flexible datacenter is no longer less costly than the critical datacenter, the critical datacenter can reassume the computing operations from the flexible datacenter.

As shown herein, by having one or more flexible datacenters powered by one or more behind-the-meter power sources available, computing operations can be managed in a dynamic manner between the critical datacenter and the flexible datacenters. The dynamic management can lower costs and, in some cases, decrease the time needed to complete time-sensitive computing operations submitted to the critical datacenter by an enterprise.

In some embodiments, one or more flexible datacenters may perform computing processes obtained through an auction process. The one or more flexible datacenters may use behind-the-meter power to acquire and perform computational operations made available via the auction process. For example, an auction process may be used to connect companies or entities requesting computational operations to be supported and performed at one or more datacenters with datacenters capable of handling the computational operations. Particularly, the auction process may involve datacenters placing bids in a competition for the various computational operations available in the auction process. For instance, the datacenter that bids to perform a computational operation at the lowest cost may win and receive the right to enter into a contract to perform the computational for the priced bid or subsequently agreed upon. As such, flexible datacenters may compete and receive the right to perform computational operations by bidding prices based on using low cost power, such as behind-the-meter power. A datacenter control system of a flexible datacenter may monitor available computational operations in multiple auctions

simultaneously to determine when to bid for computational operations based on the cost of power available and competing bids.

FIG. 1 shows a computing system 100 in accordance with one or more embodiments of the present invention. Computing system 100 may include one or more central processing units (singular “CPU” or plural “CPUs”) 105, host bridge 110, input/output (“IO”) bridge 115, graphics processing units (singular “GPU” or plural “GPUs”) 125, and/or application-specific integrated circuits (singular “ASIC” or plural “ASICs”) (not shown) disposed on one or more printed circuit boards (not shown) that are configured to perform computational operations. Each of the one or more CPUs 105, GPUs 125, or ASICs (not shown) may be a single-core (not independently illustrated) device or a multi-core (not independently illustrated) device. Multi-core devices typically include a plurality of cores (not shown) disposed on the same physical die (not shown) or a plurality of cores (not shown) disposed on multiple die (not shown) that are collectively disposed within the same mechanical package (not shown).

CPU 105 may be a general purpose computational device typically configured to execute software instructions. CPU 105 may include an interface 108 to host bridge 110, an interface 118 to system memory 120, and an interface 123 to one or more IO devices, such as, for example, one or more GPUs 125. GPU 125 may serve as a specialized computational device typically configured to perform graphics functions related to frame buffer manipulation. However, one of ordinary skill in the art will recognize that GPU 125 may be used to perform non-graphics related functions that are computationally intensive. In certain embodiments, GPU 125 may interface 123 directly with CPU 105 (and interface 118 with system memory 120 through CPU 105). In other embodiments, GPU 125 may interface 121 with host bridge 110 (and interface 116 or 118 with system memory 120 through host bridge 110 or CPU 105 depending on the application or design). In still other embodiments, GPU 125 may interface 133 with IO bridge 115 (and interface 116 or 118 with system memory 120 through host bridge 110 or CPU 105 depending on the application or design). The functionality of GPU 125 may be integrated, in whole or in part, with CPU 105.

Host bridge 110 may be an interface device configured to interface between the one or more computational devices and IO bridge 115 and, in some embodiments, system memory 120. Host bridge 110 may include an interface 108 to CPU 105, an interface 113 to IO bridge 115, for embodiments where CPU 105 does not include an interface 118 to system memory 120, an interface 116 to system memory 120, and for embodiments where CPU 105 does not include an integrated GPU 125 or an interface 123 to GPU 125, an interface 121 to GPU 125. The functionality of host bridge 110 may be integrated, in whole or in part, with CPU 105. IO bridge 115 may be an interface device configured to interface between the one or more computational devices and various IO devices (e.g., 140, 145) and IO expansion, or add-on, devices (not independently illustrated). IO bridge 115 may include an interface 113 to host bridge 110, one or more interfaces 133 to one or more IO expansion devices 135, an interface 138 to keyboard 140, an interface 143 to mouse 145, an interface 148 to one or more local storage devices 150, and an interface 153 to one or more network interface devices 155. The functionality of IO bridge 115 may be integrated, in whole or in part, with CPU 105 or host bridge 110. Each local storage device 150, if any, may be a solid-state memory device, a solid-state memory device

array, a hard disk drive, a hard disk drive array, or any other non-transitory computer readable medium. Network interface device 155 may provide one or more network interfaces including any network protocol suitable to facilitate networked communications.

Computing system 100 may include one or more network-attached storage devices 160 in addition to, or instead of, one or more local storage devices 150. Each network-attached storage device 160, if any, may be a solid-state memory device, a solid-state memory device array, a hard disk drive, a hard disk drive array, or any other non-transitory computer readable medium. Network-attached storage device 160 may or may not be collocated with computing system 100 and may be accessible to computing system 100 via one or more network interfaces provided by one or more network interface devices 155.

One of ordinary skill in the art will recognize that computing system 100 may be a conventional computing system or an application-specific computing system. In certain embodiments, an application-specific computing system may include one or more ASICs (not shown) that are configured to perform one or more functions, such as distributed computing processes or hashing, in a more efficient manner. The one or more ASICs (not shown) may interface directly with CPU 105, host bridge 110, or GPU 125 or interface through IO bridge 115. Alternatively, in other embodiments, an application-specific computing system may be reduced to only those components necessary to perform a desired function in an effort to reduce one or more of chip count, printed circuit board footprint, thermal design power, and power consumption. The one or more ASICs (not shown) may be used instead of one or more of CPU 105, host bridge 110, IO bridge 115, or GPU 125. In such systems, the one or more ASICs may incorporate sufficient functionality to perform certain network and computational functions in a minimal footprint with substantially fewer component devices.

As such, one of ordinary skill in the art will recognize that CPU 105, host bridge 110, IO bridge 115, GPU 125, or ASIC (not shown) or a subset, superset, or combination of functions or features thereof, may be integrated, distributed, or excluded, in whole or in part, based on an application, design, or form factor in accordance with one or more embodiments of the present invention. Thus, the description of computing system 100 is merely exemplary and not intended to limit the type, kind, or configuration of component devices that constitute a computing system 100 suitable for performing computing operations in accordance with one or more embodiments of the present invention.

One of ordinary skill in the art will recognize that computing system 100 may be a stand-alone, laptop, desktop, server, blade, or rack mountable system and may vary based on an application or design.

FIG. 2 shows a flexible datacenter 200 in accordance with one or more embodiments of the present invention. Flexible datacenter 200 may include a mobile container 205, a behind-the-meter power input system 210, a power distribution system 215, a climate control system (e.g., 250, 260, 270, 280, and/or 290), a datacenter control system 220, and a plurality of computing systems 100 disposed in one or more racks 240. Datacenter control system 220 may be a computing system (e.g., 100 of FIG. 1) configured to dynamically modulate power delivery to one or more computing systems 100 disposed within flexible datacenter 200 based on behind-the-meter power availability or an opera-

11

tional directive from a local station control system (not shown), a remote master control system (not shown), or a grid operator (not shown).

In certain embodiments, mobile container **205** may be a storage trailer disposed on wheels and configured for rapid deployment. In other embodiments, mobile container **205** may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical or horizontal manner (not shown). In still other embodiments, mobile container **205** may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile datacenter **200**. And in still other embodiments, flexible datacenter **200** might not include a mobile container. For example, flexible datacenter **200** may be situated within a building or another type of stationary environment.

Flexible datacenter **200** may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. Behind-the-meter power input system **210** may be configured to input power to flexible datacenter **200**. Behind-the-meter power input system **210** may include a first input (not independently illustrated) configured to receive three-phase behind-the-meter alternating current (“AC”) voltage. In certain embodiments, behind-the-meter power input system **210** may include a supervisory AC-to-AC step-down transformer (not shown) configured to step down three-phase behind-the-meter AC voltage to single-phase supervisory nominal AC voltage or a second input (not independently illustrated) configured to receive single-phase supervisory nominal AC voltage from the local station (not shown) or a metered source (not shown). Behind-the-meter power input system **210** may provide single-phase supervisory nominal AC voltage to datacenter control system **220**, which may remain powered at almost all times to control the operation of flexible datacenter **200**. The first input (not independently illustrated) or a third input (not independently illustrated) of behind-the-meter power input system **210** may direct three-phase behind-the-meter AC voltage to an operational AC-to-AC step-down transformer (not shown) configured to controllably step down three-phase behind-the-meter AC voltage to three-phase nominal AC voltage. Datacenter control system **220** may controllably enable or disable generation or provision of three-phase nominal AC voltage by the operational AC-to-AC step-down transformer (not shown).

Behind-the-meter power input system **210** may provide three phases of three-phase nominal AC voltage to power distribution system **215**. Power distribution system **215** may controllably provide a single phase of three-phase nominal AC voltage to each computing system **100** or group **240** of computing systems **100** disposed within flexible datacenter **200**. Datacenter control system **220** may controllably select which phase of three-phase nominal AC voltage that power distribution system **215** provides to each computing system **100** or group **240** of computing systems **100**. In this way, datacenter control system **220** may modulate power delivery by either ramping-up flexible datacenter **200** to fully operational status, ramping-down flexible datacenter **200** to offline status (where only datacenter control system **220** remains powered), reducing power consumption by withdrawing power delivery from, or reducing power to, one or more computing systems **100** or groups **240** of computing systems **100**, or modulating a power factor correction factor for the local station by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more computing systems **100** or groups **240** of computing systems

12

100. In some embodiments, flexible datacenter **200** may receive DC power to power computing systems **100**.

Flexible datacenter **200** may include a climate control system (e.g., **250**, **260**, **270**, **280**, **290**) configured to maintain the plurality of computing systems **100** within their operational temperature range. In certain embodiments, the climate control system may include an air intake **250**, an evaporative cooling system **270**, a fan **280**, and an air outtake **260**. In other embodiments, the climate control system may include an air intake **250**, an air conditioner or refrigerant cooling system **290**, and an air outtake **260**. In still other embodiments, the climate control system may include a computer room air conditioner system (not shown), a computer room air handler system (not shown), or an immersive cooling system (not shown). One of ordinary skill in the art will recognize that any suitable heat extraction system (not shown) configured to maintain the operation of the plurality of computing systems **100** within their operational temperature range may be used in accordance with one or more embodiments of the present invention.

Flexible datacenter **200** may include a battery system (not shown) configured to convert three-phase nominal AC voltage to nominal DC voltage and store power in a plurality of storage cells. The battery system (not shown) may include a DC-to-AC inverter configured to convert nominal DC voltage to three-phase nominal AC voltage for flexible datacenter **200** use. Alternatively, the battery system (not shown) may include a DC-to-AC inverter configured to convert nominal DC voltage to single-phase nominal AC voltage to power datacenter control system **220**.

One of ordinary skill in the art will recognize that a voltage level of three-phase behind-the-meter AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase nominal AC voltage, single-phase nominal AC voltage, and nominal DC voltage may vary based on the application or design in accordance with one or more embodiments of the present invention.

FIG. 3 shows a three-phase power distribution of a flexible datacenter **200** in accordance with one or more embodiments of the present invention. Flexible datacenter **200** may include a plurality of racks **240**, each of which may include one or more computing systems **100** disposed therein. As discussed above, the behind-the-meter power input system (**210** of FIG. 2) may provide three phases of three-phase nominal AC voltage to the power distribution system (**215** of FIG. 2). The power distribution system (**215** of FIG. 2) may controllably provide a single phase of three-phase nominal AC voltage to each computing system **100** or group **240** of computing systems **100** disposed within flexible datacenter **200**. For example, a flexible datacenter **200** may include eighteen racks **240**, each of which may include eighteen computing systems **100**. The power distribution system (**215** of FIG. 2) may control which phase of three-phase nominal AC voltage is provided to one or more computing systems **100**, a rack **240** of computing systems **100**, or a group (e.g., **310**, **320**, or **330**) of racks **240** of computing systems **100**.

In the figure, for purposes of illustration only, eighteen racks **240** are divided into a first group of six racks **310**, a second group of six racks **320**, and a third group of six racks **330**, where each rack contains eighteen computing systems **100**. The power distribution system (**215** of FIG. 2) may, for example, provide a first phase of three-phase nominal AC

13

voltage to the first group of six racks **310**, a second phase of three-phase nominal AC voltage to the second group of six racks **320**, and a third phase of three-phase nominal AC voltage to the third group of six racks **330**. If the flexible datacenter (**200** of FIG. 2) receives an operational directive from the local station (not shown) to provide power factor correction, the datacenter control system (**220** of FIG. 2) may direct the power distribution system (**215** of FIG. 2) to adjust which phase or phases of three-phase nominal AC voltage are used to provide the power factor correction required by the local station (not shown) or grid operator (not shown). One of ordinary skill in the art will recognize that, in addition to the power distribution, the load may be varied by adjusting the number of computing systems **100** operatively powered. As such, the flexible datacenter (**200** of FIG. 2) may be configured to act as a capacitive or inductive load to provide the appropriate reactance necessary to achieve the power factor correction required by the local station (not shown).

FIG. 4 shows a control distribution scheme **400** of a flexible datacenter **200** in accordance with one or more embodiments of the present invention. Datacenter control system **220** may independently, or cooperatively with one or more of local station control system **410**, remote master control system **420**, and grid operator **440**, modulate power delivery to flexible datacenter **200**. Specifically, power delivery may be dynamically adjusted based on conditions or operational directives.

Local station control system **410** may be a computing system (e.g., **100** of FIG. 1) that is configured to control various aspects of the local station (not independently illustrated) that generates power and sometimes generates unutilized behind-the-meter power. Local station control system **410** may communicate with remote master control system **420** over a networked connection **430** and with datacenter control system **220** over a networked or hardwired connection **415**. Remote master control system **420** may be a computing system (e.g., **100** of FIG. 1) that is located offsite, but connected via a network connection **425** to datacenter control system **220**, that is configured to provide supervisory or override control of flexible datacenter **200** or a fleet (not shown) of flexible datacenters **200**. Grid operator **440** may be a computing system (e.g., **100** of FIG. 1) that is configured to control various aspects of the grid (not independently illustrated) that receives power from the local station (not independently illustrated). Grid operator **440** may communicate with local station control system **410** over a networked or hardwired connection **445**.

Datacenter control system **220** may monitor unutilized behind-the-meter power availability at the local station (not independently illustrated) and determine when a datacenter ramp-up condition is met. Unutilized behind-the-meter power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, conditions where the cost for power is economically viable (e.g., low cost for power), situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution losses or costs.

The datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is

14

no operational directive from local station control system **410**, remote master control system **420**, or grid operator **440** to go offline or reduce power. As such, datacenter control system **220** may enable **435** behind-the-meter power input system **210** to provide three-phase nominal AC voltage to the power distribution system (**215** of FIG. 2) to power the plurality of computing systems (**100** of FIG. 2) or a subset thereof. Datacenter control system **220** may optionally direct one or more computing systems (**100** of FIG. 2) to perform predetermined computational operations (e.g., distributed computing processes). For example, if the one or more computing systems (**100** of FIG. 2) are configured to perform blockchain hashing operations, datacenter control system **220** may direct them to perform blockchain hashing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems (**100** of FIG. 2) may be configured to independently receive a computational directive from a network connection (not shown) to a peer-to-peer blockchain network (not shown) such as, for example, a network for a specific blockchain application, to perform predetermined computational operations.

Remote master control system **420** may specify to datacenter control system **220** what sufficient behind-the-meter power availability constitutes, or datacenter control system **220** may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter **200**. In such circumstances, datacenter control system **220** may provide power to only a subset of computing systems (**100** of FIG. 2), or operate the plurality of computing systems (**100** of FIG. 2) in a lower power mode, that is within the sufficient, but less than full, range of power that is available.

While flexible datacenter **200** is online and operational, a datacenter ramp-down condition may be met when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from local station control system **410**, remote master control system **420**, or grid operator **440**. Datacenter control system **220** may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by remote master control system **420** or datacenter control system **220** may be programmed with a predetermined preference or criteria on which to make the determination independently. An operational directive may be based on current dispatchability, forward looking forecasts for when unutilized behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the local station **410**, remote master control **420**, or grid operator **440**. For example, local station control system **410**, remote master control system **420**, or grid operator **440** may issue an operational directive to flexible datacenter **200** to go offline and power down. When the datacenter ramp-down condition is met, datacenter control system **220** may disable power delivery to the plurality of computing systems (**100** of FIG. 2). Datacenter control system **220** may disable **435** behind-the-meter power input system **210** from providing three-phase nominal AC voltage to the power distribution system (**215** of FIG. 2) to power down the plurality of computing systems (**100** of FIG. 2), while datacenter control system **220** remains powered and is capable of rebooting flexible datacenter **200** when unutilized behind-the-meter power becomes available again.

15

While flexible datacenter **200** is online and operational, changed conditions or an operational directive may cause datacenter control system **220** to modulate power consumption by flexible datacenter **200**. Datacenter control system **220** may determine, or local station control system **410**, remote master control system **420**, or grid operator **440** may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power flexible datacenter **200**. In such situations, datacenter control system **220** may take steps to reduce or stop power consumption by flexible datacenter **200** (other than that required to maintain operation of datacenter control system **220**). Alternatively, local station control system **410**, remote master control system **420**, or grid operator **440**, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, datacenter control system **220** may dynamically reduce or withdraw power delivery to one or more computing systems (**100** of FIG. 2) to meet the dictate. Datacenter control system **220** may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (**100** of FIG. 2) to reduce power consumption. Datacenter control system **220** may dynamically reduce the power consumption of one or more computing systems (**100** of FIG. 2) by reducing their operating frequency or forcing them into a lower power mode through a network directive.

One of ordinary skill in the art will recognize that datacenter control system **220** may be configured to have a number of different configurations, such as a number or type or kind of computing systems (**100** of FIG. 2) that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available unutilized behind-the-meter power availability. As such, datacenter control system **220** may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

FIG. 5 shows a control distribution of a fleet **500** of flexible datacenters **200** in accordance with one or more embodiments of the present invention. The control distribution of a flexible datacenter **200** shown and described with respect to FIG. 4 may be extended to a fleet **500** of flexible datacenters **200**. For example, a first local station (not independently illustrated), such as, for example, a wind farm (not shown), may include a first plurality **510** of flexible datacenters **200a** through **200d**, which may be collocated or distributed across the local station (not shown). A second local station (not independently illustrated), such as, for example, another wind farm or a solar farm (not shown), may include a second plurality **520** of flexible datacenters **200e** through **200h**, which may be collocated or distributed across the local station (not shown). One of ordinary skill in the art will recognize that the number of flexible datacenters **200** deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more embodiments of the present invention.

Remote master control system **420** may provide supervisory control over fleet **500** of flexible datacenters **200** in a similar manner to that shown and described with respect to FIG. 4, with the added flexibility to make high level decisions with respect to fleet **500** that may be counterintuitive to a given station. Remote master control system **420** may make decisions regarding the issuance of operational directives to a given local station based on, for example, the status of each local station where flexible datacenters **200** are deployed, the workload distributed across fleet **500**, and the

16

expected computational demand required for the expected workload. In addition, remote master control system **420** may shift workloads from a first plurality **510** of flexible datacenters **200** to a second plurality **520** of flexible datacenters **200** for any reason, including, for example, a loss of unutilized behind-the-meter power availability at one local station and the availability of unutilized behind-the-meter power at another local station.

FIG. 6 shows a flexible datacenter **200** powered by one or more wind turbines **610** in accordance with one or more embodiments of the present invention. A wind farm **600** typically includes a plurality of wind turbines **610**, each of which intermittently generates a wind-generated AC voltage. The wind-generated AC voltage may vary based on a type, kind, or configuration of farm **600**, turbine **610**, and incident wind speed. The wind-generated AC voltage is typically input into a turbine AC-to-AC step-up transformer (not shown) that is disposed within the nacelle (not independently illustrated) or at the base of the mast (not independently illustrated) of turbine **610**. The turbine AC-to-AC step up transformer (not shown) outputs three-phase wind-generated AC voltage **620**. Three-phase wind-generated AC voltage **620** produced by the plurality of wind turbines **610** is collected **625** and provided **630** to another AC-to-AC step-up transformer **640** that steps up three-phase wind-generated AC voltage **620** to three-phase grid AC voltage **650** suitable for delivery to grid **660**. Three-phase grid AC voltage **650** may be stepped down with an AC-to-AC step-down transformer **670** configured to produce three-phase local station AC voltage **680** provided to local station **690**. One of ordinary skill in the art will recognize that the actual voltage levels may vary based on the type, kind, or number of wind turbines **610**, the configuration or design of wind farm **600**, and grid **660** that it feeds into.

The output side of AC-to-AC step-up transformer **640** that connects to grid **660** may be metered and is typically subject to transmission and distribution costs. In contrast, power consumed on the input side of AC-to-AC step-up transformer **640** may be considered behind-the-meter and is typically not subject to transmission and distribution costs. As such, one or more flexible datacenters **200** may be powered by three-phase wind-generated AC voltage **620**. Specifically, in wind farm **600** applications, the three-phase behind-the-meter AC voltage used to power flexible datacenter **200** may be three-phase wind-generated AC voltage **620**. As such, flexible datacenter **200** may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

Unutilized behind-the-meter power availability may occur when there is excess local power generation. In high wind conditions, wind farm **600** may generate more power than, for example, AC-to-AC step-up transformer **640** is rated for. In such situations, wind farm **600** may have to take steps to protect its equipment from damage, which may include taking one or more turbines **610** offline or shunting their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters **200** may be used to consume power on the input side of AC-to-AC step-up transformer **640**, thereby allowing wind farm **600** to operate equipment within operating ranges while flexible datacenter **200** receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station **690** may issue an operational directive to the one or more flexible datacenters **200** or to the remote master control system (**420** of FIG. 4) to ramp-up to the desired power consumption level. When

17

the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when grid 660 cannot, for whatever reason, take the power being produced by wind farm 600. In such situations, wind farm 600 may have to take one or more turbines 610 offline or shunt their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 640, thereby allowing wind farm 600 to either produce power to grid 660 at a lower level or shut down transformer 640 entirely while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 or the grid operator (not independently illustrated) of grid 660 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when wind farm 600 is selling power to grid 660 at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price wind farm 600 would have to pay to grid 660 to offload their generated power. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to produce and obtain the production tax credit, but sell less power to grid 660 at the negative price. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when wind farm 600 is selling power to grid 660 at a negative price because grid 660 is oversupplied or is instructed to stand down and stop producing altogether. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 660. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to stop producing power to grid 660, but making productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of the local station 690 or the grid operator (not independently illustrated) of grid 660 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the coopera-

18

tive action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when wind farm 600 is producing power to grid 660 that is unstable, out of phase, or at the wrong frequency, or grid 660 is already unstable, out of phase, or at the wrong frequency for whatever reason. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 660. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to stop producing power to grid 660, but make productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Further examples of unutilized behind-the-meter power availability is when wind farm 600 experiences low wind conditions that make it not economically feasible to power up certain components, such as, for example, the local station (not independently illustrated), but there may be sufficient behind-the-meter power availability to power one or more flexible datacenters 200. Similarly, unutilized behind-the-meter power availability may occur when wind farm 600 is starting up, or testing, one or more turbines 610. Turbines 610 are frequently offline for installation, maintenance, and service and must be tested prior to coming online as part of the array. One or more flexible datacenters 200 may be powered by one or more turbines 610 that are offline from farm 600. The above-noted examples of when unutilized behind-the-meter power is available are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as unutilized behind-the-meter power availability. Unutilized behind-the-meter power availability may occur anytime there is power available and accessible behind-the-meter that is not subject to transmission and distribution costs and there is an economic advantage to using it.

One of ordinary skill in the art will recognize that wind farm 600 and wind turbine 610 may vary based on an application or design in accordance with one or more embodiments of the present invention.

FIG. 7 shows a flexible datacenter 200 powered by one or more solar panels 710 in accordance with one or more embodiments of the present invention. A solar farm 700 typically includes a plurality of solar panels 710, each of which intermittently generates a solar-generated DC voltage 720. Solar-generated DC voltage 720 may vary based on a type, kind, or configuration of farm 700, panel 710, and incident sunlight. Solar-generated DC voltage 720 produced by the plurality of solar panels 710 is collected 725 and provided 730 to a DC-to-AC inverter 740 that converts solar-generated DC voltage into three-phase solar-generated AC voltage 750. Three-phase solar-generated AC voltage 750 is provided to an AC-to-AC step-up transformer 760

19

that steps up three-phase solar-generated AC voltage to three-phase grid AC voltage 790. Three-phase grid AC voltage 790 may be stepped down with an AC-to-AC step-down transformer 785 configured to produce three-phase local station AC voltage 777 provided to local station 775. One of ordinary skill in the art will recognize that the actual voltage levels may vary based on the type, kind, or number of solar panels 710, the configuration or design of solar farm 700, and grid 790 that it feeds into. In some embodiments, the solar farm 700 may provide DC power directly to flexible datacenter 200 without a conversion to AC via the DC-to-AC inverter 740.

The output side of AC-to-AC step-up transformer 760 that connects to grid 790 may be metered and is typically subject to transmission and distribution costs. In contrast, power consumed on the input side of AC-to-AC step-up transformer 760 may be considered behind-the-meter and is typically not subject to transmission and distribution costs. As such, one or more flexible datacenters 200 may be powered by three-phase solar-generated AC voltage 750. Specifically, in solar farm 700 applications, the three-phase behind-the-meter AC voltage used to power flexible datacenter 200 may be three-phase solar-generated AC voltage 750. As such, flexible datacenter 200 may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

Unutilized behind-the-meter power availability may occur when there is excess local power generation. In high incident sunlight situations, solar farm 700 may generate more power than, for example, AC-to-AC step-up transformer 760 is rated for. In such situations, solar farm 700 may have to take steps to protect its equipment from damage, which may include taking one or more panels 710 offline or shunting their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 760, thereby allowing solar farm 700 to operate equipment within operating ranges while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when grid 790 cannot, for whatever reason, take the power being produced by solar farm 700. In such situations, solar farm 700 may have to take one or more panels 710 offline or shunt their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 760, thereby allowing solar farm 700 to either produce power to grid 790 at a lower level or shut down transformer 760 entirely while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 or the grid operator (not independently illustrated) of grid 790 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420

20

of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when solar farm 700 is selling power to grid 790 at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price solar farm 700 would have to pay to grid 790 to offload their generated power. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to produce and obtain the production tax credit, but sell less power to grid 790 at the negative price. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenter 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when solar farm 700 is selling power to grid 790 at a negative price because grid 790 is oversupplied or is instructed to stand down and stop producing altogether. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 790. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to stop producing power to grid 790, but making productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of the local station 775 or the grid operator (not independently illustrated) of grid 790 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when solar farm 700 is producing power to grid 790 that is unstable, out of phase, or at the wrong frequency, or grid 790 is already unstable, out of phase, or at the wrong frequency for whatever reason. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 790. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to stop producing power to grid 790, but make productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the

21

operational directive requires the cooperative action of multiple flexible datacenters **200**, the remote master control system (**420** of FIG. **4**) may determine how to power each individual flexible datacenter **200** in accordance with the operational directive or provide an override to each flexible datacenter **200**.

Further examples of unutilized behind-the-meter power availability is when solar farm **700** experiences intermittent cloud cover such that it is not economically feasible to power up certain components, such as, for example local station **775**, but there may be sufficient behind-the-meter power availability to power one or more flexible datacenters **200**. Similarly, unutilized behind-the-meter power availability may occur when solar farm **700** is starting up, or testing, one or more panels **710**. Panels **710** are frequently offline for installation, maintenance, and service and must be tested prior to coming online as part of the array. One or more flexible datacenters **200** may be powered by one or more panels **710** that are offline from farm **700**. The above-noted examples of when unutilized behind-the-meter power is available are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as unutilized behind-the-meter power availability. Behind-the-meter power availability may occur anytime there is power available and accessible behind-the-meter that is not subject to transmission and distribution costs and there is an economic advantage to using it.

One of ordinary skill in the art will recognize that solar farm **700** and solar panel **710** may vary based on an application or design in accordance with one or more embodiments of the present invention.

FIG. **8** shows a flexible datacenter **200** powered by flare gas **800** in accordance with one or more embodiments of the present invention. Flare gas **800** is combustible gas produced as a product or by-product of petroleum refineries, chemical plants, natural gas processing plants, oil and gas drilling rigs, and oil and gas production facilities. Flare gas **800** is typically burned off through a flare stack (not shown) or vented into the air. In one or more embodiments of the present invention, flare gas **800** may be diverted **812** to a gas-powered generator that produces three-phase gas-generated AC voltage **822**. This power may be considered behind-the-meter and is not subject to transmission and distribution costs. As such, one or more flexible datacenters **200** may be powered by three-phase gas-generated AC voltage. Specifically, the three-phase behind-the-meter AC voltage used to power flexible datacenter **200** may be three-phase gas-generated AC voltage **822**. Accordingly, flexible datacenter **200** may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

FIG. **9A** shows a method of dynamic power delivery to a flexible datacenter (**200** of FIG. **2**) using behind-the-meter power **900** in accordance with one or more embodiments of the present invention. In step **910**, the datacenter control system (**220** of FIG. **4**), or the remote master control system (**420** of FIG. **4**), may monitor behind-the-meter power availability. In certain embodiments, monitoring may include receiving information or an operational directive from the local station control system (**410** of FIG. **4**) or the grid operator (**440** of FIG. **4**) corresponding to behind-the-meter power availability.

In step **920**, the datacenter control system (**220** of FIG. **4**), or the remote master control system (**420** of FIG. **4**), may determine when a datacenter ramp-up condition is met. In certain embodiments, the datacenter ramp-up condition may

22

be met when there is sufficient behind-the-meter power availability and there is no operational directive from the local station to go offline or reduce power. In step **930**, the datacenter control system (**220** of FIG. **4**) may enable behind-the-meter power delivery to one or more computing systems (**100** of FIG. **2**). In step **940**, once ramped-up, the datacenter control system (**220** of FIG. **4**) or the remote master control system (**420** of FIG. **4**) may direct one or more computing systems (**100** of FIG. **2**) to perform predetermined computational operations. In certain embodiments, the predetermined computational operations may include the execution of one or more distributed computing processes, parallel processes, and/or hashing functions, among other types of processes.

While operational, the datacenter control system (**220** of FIG. **4**), or the remote master control system (**420** of FIG. **4**), may receive an operational directive to modulate power consumption. In certain embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system (**220** of FIG. **4**) or the remote master control system (**420** of FIG. **4**) may dynamically reduce power delivery to one or more computing systems (**100** of FIG. **2**) or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system (**220** of FIG. **4**) or the remote master control system (**420** of FIG. **4**) may dynamically adjust power delivery to one or more computing systems (**100** of FIG. **2**) to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system (**220** of FIG. **4**) may disable power delivery to one or more computing systems (**100** of FIG. **2**).

As such, FIG. **9B** shows a method of dynamic power delivery to a flexible datacenter (**200** of FIG. **2**) using behind-the-meter power **950** in accordance with one or more embodiments of the present invention. In step **960**, the datacenter control system (**220** of FIG. **4**), or the remote master control system (**420** of FIG. **4**), may monitor behind-the-meter power availability. In certain embodiments, monitoring may include receiving information or an operational directive from the local station control system (**410** of FIG. **4**) or the grid operator (**440** of FIG. **4**) corresponding to behind-the-meter power availability.

In step **970**, the datacenter control system (**220** of FIG. **4**), or the remote master control system (**420** of FIG. **4**), may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met when there is insufficient behind-the-meter power availability or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the local station to go offline or reduce power. In step **980**, the datacenter control system (**220** of FIG. **4**) may disable behind-the-meter power delivery to one or more computing systems (**100** of FIG. **2**). In step **990**, once ramped-down, the datacenter control system (**220** of FIG. **4**) remains powered and in communication with the remote master control system (**420** of FIG. **4**) so that it may dynamically power the flexible datacenter (**200** of FIG. **2**) when conditions change.

One of ordinary skill in the art will recognize that a datacenter control system (**220** of FIG. **4**) may dynamically modulate power delivery to one or more computing systems (**100** of FIG. **2**) of a flexible datacenter (**200** of FIG. **2**) based on behind-the-meter power availability or an operational

23

directive. The flexible datacenter (200 of FIG. 2) may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter (200 of FIG. 2) may have a blackout state, where all power consumption, including that of the datacenter control system (220 of FIG. 4) is halted. However, once the flexible datacenter (200 of FIG. 2) enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system (220 of FIG. 4). Local station conditions or operational directives may cause flexible datacenter (200 of FIG. 2) to ramp-up, reduce power consumption, change power factor, or ramp-down.

FIG. 10 illustrates a system for managing queue distribution among a critical datacenter and behind-the-meter flexible datacenters in accordance with one or more embodiments of the present invention. The system 1000 includes a flexible datacenter 200, a critical datacenter 1004, a communication link 1006, a queue system 1008, a communication link 1010, communication links 1012a, 1012b, and a remote master control system 420. The system 1000 represents an example configuration scheme for a system that can distribute computing operations using a queue system 1008 between the critical datacenter 1004 and one or more flexible datacenters (e.g., the flexible datacenter 200). In other examples, the system 1000 may include more or fewer components in other potential configurations. For instance, the system 1000 may not include the queue system 1008 or may include routes that bypass the queue system 1008.

The system 1000 may be configured to manage computational operations requested to be performed by enterprises or other entities. Computational operations may include various tasks that can be performed or generally supported by one or more computing systems within a datacenter. The parameters of each set of computational operations submitted by an enterprise may differ. For instance, the amount of computational resources (e.g., number of computing systems), the degree of difficulty, the duration and degree of support required, etc., may vary for each set of computational operations. As such, the system 1000 may use the queue system 1008 to organize incoming computational operations requests to enable efficient distribution to the flexible datacenter 200 and the critical datacenter 1004. The system 1000 may use the queue system 1008 to organize sets of computational operations thereby increasing the speed of distribution and performance of the different computational operations requested by various enterprises as compared to a system without a queue system 1008.

In some examples, the queue system 1008 may enable the system 1000 to efficiently utilize the flexible datacenter 200 to perform some sets of computational operations in a manner that can reduce costs or time required to complete the sets. In particular, one or more components within the system 1000, such as the control systems 220, 420, or 1022, may be configured to identify situations that may arise where using the flexible datacenter 200 can reduce costs or increase productivity of the system 1000, as compared to using the critical datacenter 1004 for computational operations. For example, a component within the system 1000, such as the control systems 220, 420, or 1022, may identify when using behind-the-meter power to power the computing systems 100 within the flexible datacenter 200 is at a lower cost compared to using the computing systems 1020 within the critical datacenter 1004, which are powered by grid power. Additionally, a component in the system 1000, control systems 220, 420, or 1022, may be configured to determine situations when offloading computational opera-

24

tions from the critical datacenter 1004 indirectly (i.e., via the queue system 1008) or directly (i.e., bypassing the queue system 1008) to the flexible datacenter 200 can increase the performance allotted to the computational operations requested by an enterprise (e.g., reduce the time required to complete time-sensitive computational operations).

Within system 1000, the flexible datacenter 200 may represent one or more flexible datacenters capable of offering computational processing and other computing resources using behind-the-meter power from behind-the-meter sources, such as illustrated in FIGS. 6, 7, and 8. As shown in FIG. 10, the flexible datacenter 200 may include a behind-the-meter power input system 215 that is connected to a behind-the-meter power source, a power distribution system 215, computing systems 100, and a datacenter control system 220, and may take the form of a mobile container or another configuration. The flexible datacenter 200 may additionally be connected to other power sources, such as other behind-the-meter power sources, the power grid, and/or an energy storage system. Additionally, the flexible datacenter 200 may include other components not shown in FIG. 10, such as a climate control system.

The location of the flexible datacenter 200 relative to the critical datacenter 1004 can vary within embodiments. In some examples, the flexible datacenter 200 may be collocated with critical datacenter 1004. For instance, one or more flexible datacenters 200 may be positioned in the same general location as the critical datacenter 1004 or even share a building with the critical datacenter 1004. In other examples, the flexible datacenter 200 and the critical datacenter 1004 are not collocated. Particularly, one or more flexible datacenters 200 within the system 1000 can have a different location from the critical datacenter 1004. In further examples, one or more flexible datacenters 200 can share a location with the critical datacenter 1004 while other flexible datacenters 200 can have a location away from the critical datacenter 1004.

In order to provide computing resources to perform or support computational operations, the flexible datacenter 200 may be deployed near or otherwise connected to one or more sources of behind-the-meter power generation. For instance, one or more flexible datacenters 200 may be connected behind-the-meter to the wind farm 600, the solar farm 700, and/or other potentially intermittent power generation sources. As such, the behind-the-meter power input system 210 may be configured to receive behind-the-meter power from one or more sources and input the power to the flexible datacenter 200. For example, the behind-the-meter power input system 210 may provide three-phase nominal AC voltage to the power distribution system 215. The power distribution system 215 may controllably provide a single phase of three-phase nominal AC voltage to one or more computing systems 100 of flexible datacenter 200. For instance, power distribution system 215 may distribute power to the computing systems 100 individually or according to groups of computing systems. The computing systems 100 may then use the power received from the behind-the-meter sources to provide processing/computing abilities, networking, storage, and other resources. In some examples, the computing systems 100 may include one or more ASIC computing systems, GPU computing systems, and/or CPU computing systems.

In some examples, power received at the flexible datacenter 200 may actively switch between different behind-the-meter sources. For example, the flexible datacenter 200 may actively switch from receiving power from either or both the wind farm 600 and the solar farm 700 (or other

25

types of sources). A control system associated with the flexible datacenter **200** (e.g., the datacenter control system **220**) or associated with the system **1000** (e.g., remote master control system **420**) generally may monitor various input signals, such as, but not limited to, the price for power, availability of power, computing analysis, and order from an operator, etc., to determine which sources to receive power from at a given time. In some situations, the control system may determine that no source is currently a viable option for supplying power to the flexible datacenter **200**. Other sources of behind-the-meter power or grid power can also be used to power the flexible datacenter **200** within examples. For example, the flexible datacenter **200** may receive grid power from the local station where it is cited.

In some examples, the datacenter control system **220** may monitor activity of the computing systems **100** within the flexible datacenter **200** and use the activity to determine when to obtain computational operations from the queue system **1008**. The datacenter control system **220** may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems **100** to perform. The various factors may include power availability at the flexible datacenter **200**, availability of the computing systems **100**, type of computational operations available, estimated cost to perform the computational operations at the flexible datacenter **200**, cost for power, cost for power relative to cost for grid power, and instructions from other components within the system **1000**, among others. The datacenter control system **220** may analyze one or more of the factors when determining whether to obtain a new set of computational operations for the computing systems **100** to perform. In such a configuration, the datacenter control system **220** manages the activity of the flexible datacenter **200**, including determining when to acquire new sets of computational operations when capacity among the computing systems **100** permit.

In other examples, a component (e.g., the remote master control system **420**) within the system **1000** may assign or distribute one or more sets of computational operations organized by the queue system **1008** to the flexible datacenter **200**. For example, the remote master control system **420** may manage the queue system **1008**, including the distribution of computational operations organized by the queue system **1008** to the flexible datacenter **1002** and the critical datacenter **1004**.

The system **1000** also includes the critical datacenter **1004**, which represents one or more datacenters assigned to provide computational resources to fulfill critical operations. Particularly, the critical datacenter **1004** may receive one or more assignments to support computational operations from an enterprise. In some examples, the critical datacenter **1004** may receive sets of computational operations directly from the enterprise or the remote master control system **420**. The critical datacenter **1004** may also receive sets of computational operations organized by the queue system **1008**. As such, to warrant that critical operations are supported, the critical datacenter **1004** is preferably connected to a power grid to ensure that reliable (i.e., non-intermittent) power is available.

The critical datacenter **1004** may include a power input system **1016**, a power distribution system **1018**, a critical datacenter control system **1022**, and computing systems **1020**. The power input system **1016** may be configured to receive power from a power grid and distribute the power to the computing systems **1020** via the power distribution system **1018**. In some embodiments, the critical datacenter

26

control system **1022** can manage the assignment and support of computational operations received from enterprises, including the distribution of computational operations among the flexible datacenter **200** and the critical datacenter **1004**. This is further described below with respect to remote master control system **420**, and management operations described with respect to remote master control system **420** may alternatively or additionally be handled by critical datacenter control system **1022**.

Similar to the flexible datacenter, the critical datacenter **1004** may access and obtain sets of computational operations organized by the queue system **1008**. The critical datacenter control system **1022** may monitor activity of the computing systems **1020** and obtain computational operations to perform from the queue system **1008**. The critical datacenter control system **1022** may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems **1020** to perform. Various factors may include power availability at the critical datacenter **1004**, power availability at the flexible datacenter **200**, availability of the computing systems **1020**, type of computational operations available, cost for power from the grid, estimated cost for the critical datacenter **1004** to perform the set computational operations, and instructions from other components within the system **1000**, among others. In other examples, a component (e.g., the remote master control system **420**) within the system **1000** may assign or distribute one or more sets of computational operations organized by the queue system **1008** to the critical datacenter **1004**.

The communication link **1006** represents one or more links that may serve to connect the flexible datacenter **200**, the critical datacenter **1004**, and other components within the system **1000** (e.g., the remote master control system **420**, the queue system **1008**—connections not shown). In particular, the communication link **1006** may enable direct or indirect communication between the flexible datacenter **200** and the critical datacenter **1004**. The type of communication link **1006** may depend on the locations of the flexible datacenter **200** and the critical datacenter **1004**. Within embodiments, different types of communication links can be used, including but not limited to WAN connectivity, cloud-based connectivity, and wired and wireless communication links.

The queue system **1008** represents an abstract data type capable of organizing computational operation requests received from enterprises. As each request for computational operations are received, the queue system **1008** may organize the request in some manner for subsequent distribution to a datacenter.

Different types of queues can make up the queue system **1008** within embodiments. The queue system **1008** may be a centralized queue that organizes all requests for computational operations. As a centralized queue, all incoming requests for computational operations may be organized by the centralized queue.

In other examples, the queue system **1008** may be distributed consisting of multiple queue subsystems. In the distributed configuration, the queue system **1008** may use multiple queue subsystems to organize different sets of computational operations. Each queue subsystem may be used to organize computational operations based on various factors, such as according to deadlines for completing each set of computational operations, locations of enterprises submitting the computational operations, economic value associated with the completion of computational operations, and quantity of computing resources required to perform

each set of computational operations. For instance, a first queue subsystem may organize sets of non-intensive computational operations and a second queue subsystem may organize sets of intensive computational operations.

Within the system 1000, the queue system 1008 is shown connected to the remote master control system 420 via the communication link 1010. In addition, the queue system 1008 is also shown connected to the flexible datacenter via the communication 1012a and to the critical datacenter 1004 via the communication link 1012b. The communication links 1010, 1012a, 1012b may be similar to the communication link 1006 and can be various types of communication links within examples.

The organizational design of the queue system 1008 may vary within examples. In some examples, the queue system 1008 may organize indications (e.g., tags, pointers to) to sets of computational operations requested by various enterprises. The queue system 1008 may operate as a First-In-First-Out (FIFO) data structure. In a FIFO data structure, the first element added to the queue will be the first one to be removed. As such, the queue system 1008 may include one or more queues that operate using the FIFO data structure.

In some examples, one or more queues within the queue system 1008 may use other configurations of queues, including rules to rank or organize queues in a particular manner that can prioritize some sets of computational operations over others. The rules may include one or more of an estimated cost and/or revenue to perform each set of computational operations, an importance assigned to each set of computational operations, and deadlines for initiating or completing each set of computational operations, among others.

The queue system 1008 may include a computing system configured to organize and maintain queues within the queue system 1008. In another example, one or more other components of the system 1000 may maintain and support queues within the queue system 1008. For instance, the remote master control system 420 may maintain and support the queue system 1008. In other examples, multiple components may maintain and support the queue system 1008 in a distributed manner, such as a blockchain configuration.

In some examples, the queue system 1008 may include queue subsystems located at each datacenter. This way, each datacenter (e.g., via a datacenter control system) may organize computational operations obtained at the datacenter until computing systems are able to start executing the computational operations.

The remote master control system 420 represents a component within the system 1000 that, in some embodiments, can manage the assignment and support of computational operations received from enterprises, including the distribution of computational operations among the flexible datacenter 200 and the critical datacenter 1004. As shown in FIG. 10, the remote master control system 420 may connect to the flexible datacenter 200 via communication link 425 and the critical datacenter 1004 via communication link 1002. Alternatively or additionally, remote master control system 420 may connect to flexible datacenter 200 and critical datacenter 1004 via communication link 1006 (not shown) or alternative communication links.

In some embodiments, remote master control system 420 may serve as an intermediary that facilitates all communication between flexible datacenter 200 and critical datacenter 1004. Particularly, critical datacenter 1004 or flexible datacenter 200 might need to transmit communications to remote master control system 420 in order to communicate with the other datacenter. As also shown, the remote master

control system 420 may connect to the queue system 1008 via the communication link 1010. Computational operations may be distributed between the queue system 1008 and the remote master control system 420 via the communication link 1010.

The remote master control system 420 may assist with management of operations assigned to one or both of the flexible datacenter 200 and the critical datacenter 1004. For instance, the remote master control system 420 may be configured to monitor input signals from behind-the-meter sources in order to identify situations where utilizing the flexible datacenter 200 can reduce costs or increase efficiency of the system 1000. For instance, the remote master control system 420 may determine when flexible datacenter 200 could use power from one or more behind-the-meter power sources to advantageously supplement the computing resources offered by the critical datacenter 1004.

In addition, the remote master control system 420 may manage the queue system 1008, including providing resources to support queues within the queue system 1008. The remote master control system 420 may also manage the distribution of computational resources from the queue system 1008 to the flexible datacenter 200 and the critical datacenter 1004. The remote master control system 420 may communicate with the datacenter control system 220 in the flexible datacenter 200 via the communication link 425 and the critical datacenter control system 1022 in the critical datacenter 1004 via the communication link 1012b. The communication may include checking whether either datacenter can receive sets of operations. This way, the remote master control system 420 may distribute sets of computational operations organized by the queue system 1008, efficiently increasing the rate which computational operations are completed.

As an example, the remote master control system 420 (or another component within the system 1000) may determine when power from a behind-the-meter source is being sold at a negative price back to the grid. As another example, the remote master control system 420 may monitor power system conditions and issue operational directives to the flexible datacenter 200. Operational directives may include, but are not limited to, a local station directive, a remote master control directive, a grid directive, a dispatchability directive, a forecast directive, a workload directive based on actual behind-the-meter power availability or projected behind-the-meter power availability. Power system conditions, which may additionally or alternatively be monitored by one or more of the control systems 220, 420, and/or 1020 may include, but are not limited to, excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation subject to economic curtailment, local power generation subject to reliability curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation. As another example, remote master control system 420 (or critical datacenter control system 1022) may monitor the types of computational operations requested of the critical datacenter 1004 and make determinations alone or in conjunction with other control systems, power system conditions, and/or operational directives to decide when or how to offload computational operations to a flexible datacenter 200.

29

As a result, the remote master control system **420** may offload some or all of the computational operations assigned to the critical datacenter **1004** to the flexible datacenter **200**. This way, flexible datacenter **200** can reduce overall computational costs by using the behind-the-meter power to provide computational resources to assist critical datacenter **1004**. The remote master control system **420** may use the queue system **1008** to temporarily store and organize the offloaded computational operations until a flexible datacenter (e.g., the flexible datacenter **200**) is available to perform them. The flexible datacenter **200** consumes behind-the-meter power without transmission or distribution costs, which lowers the costs associated with performing computational operations originally assigned to critical datacenter **1004**.

In further examples, remote master control system **420** may identify other situations that may benefit from using one or more flexible datacenters (e.g., flexible datacenter **200**) to supplement or replace computational resources provided by critical datacenter **1004**.

In some examples, remote master control system **420** may facilitate communication among components within system **1000** using communication links **425**, **1002**, **1006**, **1010**, **1012a**, and/or **1012b**. The communications may include computation requests from components within system **1000**. In one embodiment, the remote master control system **420** may identify a computational operation to be performed at a critical datacenter **1004**. The computational operation may be identified by querying the critical datacenter **1004** or by receiving a request from the critical datacenter **1004**. Information regarding active or requested computational operations at the critical datacenter **1004** may be considered as part of the identification process. The communications may also include a variety of other information, such as an indication of a current workload at the critical datacenter **1004**, a current status of operation at critical datacenter **1004** (e.g., a report indicating current capacity available and power consumption at critical datacenter **1004**). Upon receiving the information, the remote master control system **420** may determine whether to route the computational operations to the flexible datacenter **200**.

The determination process may involve considering various factors, including power availability and associated costs from the power grid and behind-the-meter sources, availability of flexible datacenter **200**, and type and deadlines associated with assigned computational operations, among others. In some situations, remote master control system **420** may then send the computational operation to flexible datacenter **200** (e.g., via communication link **1006**). In these situations, remote master control system **420** may determine that utilizing the flexible datacenter **200** could enhance the operation of system **1000** overall (i.e. improving profitability or timely performance). Particularly, using the flexible datacenter **200** may reduce costs and increase efficiency of system **1000**. The flexible datacenter **200** may also help reduce the amount of unutilized or under-utilized power being produced by one or more behind-the-meter sources.

In some examples, the remote master control system **420** may reassign computational operations from critical datacenter **1004** over to the flexible datacenter **200** for the flexible datacenter **200** to support or complete. For instance, the remote master control system **420** may determine that using the flexible datacenter **200** is more cost efficient than only using critical datacenter **1004**. As such, the remote master control system **420** may facilitate a direct transfer of responsibility for the computational operations from the critical datacenter **1004** to the flexible datacenter **200**. Alter-

30

natively, the remote master control system **420** may use the queue system **1008** to facilitate an indirect transfer of computational operations from the critical datacenter **1004** to the flexible datacenter **200**. Particularly, the remote master control system **420** may transfer the offloaded computational operations from the critical datacenter into the queue system **1008** until a flexible datacenter **200** is able to perform the computational operations. The flexible datacenter **200** may access and obtain the offloaded computational operations from the queue system **1008** or may be assigned the computational operations from the queue system **1008** by the remote master control system **420** or another component within the system **1000**.

In further examples, the remote master control system **420** may determine that the flexible datacenter **200** is available to support and provide computing resources to new computational operations received from an enterprise. This way, the remote master control system **420** may route the new computational operations directly to the flexible datacenter **200** or indirectly via use of the queue system **1008** without impacting the workload on the critical datacenter **1004**.

When determining whether to route a computational operation to the flexible datacenter **200**, the remote master control system **420** may be configured to consider different factors, such as the availability of the flexible datacenter **200** and availability of behind-the-meter power. In some situations, the remote master control system **420** or another component within the system **1000** (e.g., datacenter control system **220**) may determine that the flexible datacenter **200** might not have enough computing systems **100** available to satisfy the computational operation. As a result, the remote master control system **420** may refrain from sending the computational operation to flexible datacenter **200**. The remote master control system **420** may then transmit an indication that the flexible datacenter **200** is unavailable back to the critical datacenter **1004**. In some examples, the remote master control system **420** may utilize the queue system **1008** to organize the computational requests until a flexible datacenter or the critical datacenter **1004** is available.

In some examples, the remote master control system **420** may further analyze the workloads of other flexible datacenters to identify a flexible datacenter that is capable of handling the computational operation. Upon identifying an available flexible datacenter, the remote master control system **420** may transmit the computational operation to that flexible datacenter instead. In further examples, the remote master control system **420** may divide operations associated with one or more identified computational operation among multiple flexible datacenters.

In some examples, the remote master control system **420** may determine whether to route a computational operation to the flexible datacenter **200** based on the availability of between-the-meter power for the flexible datacenter **200**. Additionally or alternatively, the remote master control system **420**, the flexible datacenter control system **220**, or another computing device may monitor one or more other power system operation conditions to make the determination. The remote master control system **420** may also determine whether a datacenter ramp-up condition is met when determining whether to route a computational operation to the flexible datacenter **200**. For instance, the remote master control system **420** may check whether the flexible datacenter **200** is ramped-up to a fully online status, ramped-down to a fully offline status, or in another state (e.g., acting as a load balancer). As such, the remote master control

31

system 420 may determine whether to route a computation request to the flexible datacenter 200 based on the status of the flexible datacenter 200.

As previously discussed, the system 1100 may include a flexible datacenter control system 220, which may be configured to modulate power delivery to computing systems 100 of flexible datacenter 200. For example, the flexible datacenter control system 220 may modulate power delivery to the computing systems 100 based on a threshold level of unutilized behind-the-meter power availability or some other monitored power system condition. In some instances, the flexible datacenter control system 220 may be configured to modulate power delivery to computing systems 100 by selectively enabling or disabling a subset of computing systems 100.

The flexible datacenter control system 220 may alternatively or additionally be configured to modulate power delivery to the computing systems 100 based on an operational directive. For instance, the flexible datacenter control system 220 or another system may receive an operational directive from a user interface to modulate the power delivery to computing systems 100. As discussed above, the operational directive may be a local station directive, a remote master control directive, a grid directive, a dispatchability directive, or a forecast directive. In some instances, the operational directive may also include a workload directive based on a threshold level actual behind-the-meter power availability or a threshold level of projected behind-the-meter power availability.

FIG. 11 illustrates a system for managing queue distribution among a critical datacenter and a plurality of behind-the-meter flexible datacenters in accordance with one or more embodiments of the present invention. The system 1100 is similar to the schemes illustrated in FIG. 5, with the addition of the critical datacenter 1004, the queue subsystem 1008a, the queue subsystem 1008b, and communication links 1002, 1006a, 1006b, 1014a, and 1014b. Local stations 410a and 410b, and other control paths not required for illustrative purposes, are removed for clarity. Components and aspects illustrated and/or described in FIG. 10 that are similar or the same as components or aspects illustrated and/or described in FIG. 5 should be considered to have the same characteristics as previously illustrated and/or described.

The system 1100 may operate similarly to the system 1000 shown in FIG. 10. Similarly labeled components in FIG. 11 may have the same characteristics and/or capabilities as described with respect to FIG. 10. The system 1100, however, is shown configured with the queue system 1008 arranged in a distributed queue subsystem format (e.g., a queue subsystem 1008a and a queue subsystem 1008b). Although queue subsystems 1008a, 1008b are shown physically separate, the queue subsystems 1008a, 1008b may be collocated and/or supported by the same component of the system 1100.

The system 1100 includes the queue subsystem 1008a connected to the remote master control system 420 via the communication link 1014a and to flexible datacenters 200a, 200b, 200c, 200d and the critical datacenter 1004 via the communication link 1006a. In such a configuration, the remote master control system 420, the flexible datacenters 200a, 200b, 200c, 200d and the critical datacenter 1004 may communicate with the queue subsystem 1008a. The communication may involve obtaining computational operations organized by the queue subsystem 1008a.

In some examples, the remote master control system 420 may distribute computational operations organized by the

32

queue subsystem 1008a to the flexible datacenters 200a, 200b, 200c, 200d and the critical datacenter 1004. In some examples, the remote master control system 420 may distribute computational operations organized by the queue subsystem 1008a to the flexible datacenters 200e, 200f, 200g, 200h, which are shown not directly connected to the queue subsystem 1008a.

In some embodiments, the flexible datacenters 200a, 200b, 200c, 200d and the critical datacenter 1004 may communicate directly with the queue subsystem 1008a to obtain computational operations to perform. This way, control systems at each datacenter may monitor and balance the workload supported by their computing systems.

The system 1100 also includes the queue subsystem 1008b connected to the remote master control system 420 via the communication link 1014b and to flexible datacenters 200e, 200f, 200g, 200h and the critical datacenter 1004 via the communication link 1006b. In such a configuration, the remote master control system 420, the flexible datacenters 200e, 200f, 200g, 200h and the critical datacenter 1004 may communicate with the queue subsystem 1008b. The communication may involve obtaining computational operations organized by the queue subsystem 1008b.

In some examples, the remote master control system 420 may distribute computational operations organized by the queue subsystem 1008b to the flexible datacenters 200e, 200f, 200g, 200h and the critical datacenter 1004. In some examples, the remote master control system 420 may distribute computational operations organized by the queue subsystem 1008b to the flexible datacenters 200a, 200b, 200c, 200d, which are shown not directly connected to the queue subsystem 1008b.

In some embodiments, the flexible datacenters 200e, 200f, 200g, 200h and the critical datacenter 1004 may communicate directly with the queue subsystem 1008b to obtain computational operations to perform. This way, control systems at each datacenter may monitor and balance the workload supported by their computing systems.

In some embodiments, the remote master control system 420 may be configured to determine whether to route a computational operation to a particular flexible datacenter (e.g., flexible datacenter 200a) from among multiple flexible datacenters. The determination process may involve initially determining whether to route the computational operation to a flexible datacenter and then further selecting a specific flexible datacenter to route the computational operation to. The remote master control system 420 or another component (e.g., one or more flexible datacenter control systems 220) may be configured to determine a cost of execution of the computing instructions by computing systems at the specific flexible datacenter. Particularly, each flexible datacenter may be capable of providing computing resources at different costs based on various factors, such as the locations of the flexible datacenters 200 and the availability of behind-the-meter power to each flexible datacenter (i.e., the flexible datacenters 200 may connect to different behind-the-meter power sources). As such, the remote master control system 420 may be configured to consider the different factors to select a specific flexible datacenter to use to fulfill a computational operation. In some examples, the remote master control system 420 may use the queue subsystems 1008a, 1008b to temporarily organize sets of computational operations until a flexible datacenter or the critical datacenter 1004 is available to perform each set.

FIG. 12 illustrates a method for managing queue distribution between a critical datacenter and a flexible datacenter in accordance with one or more embodiments of the present

33

invention. The method serves an example and may include other steps within other examples. At step **1202**, the method involves identifying, using a queue system, a computational operation to be performed. For instance, a component within the system **1000** may identify a computational operation to be performed by analyzing the queue system **1008**. The component may be the remote master control system **420**, the datacenter control system **220**, the critical datacenter control system **1022**, and/or a computing system at the queue system **1008** in some examples.

Identifying the computational operation can include examining various types of information, such as a request for processing, networking, or storage capabilities or a request to offload some work from the critical datacenter. In some instances, the computational operation may be identified in association with an incoming computational operation request received from an outside enterprise. In some examples, the computational operation may be identified based on the organization of the queue system **1008**. For instance, the computational operation may be the next operation to be selected based on a FIFO format of the queue system **1008**.

At step **1204**, the method involves determining whether to route the computational operation to a flexible datacenter. Different components may be configured to determine whether to route the computational operation to a flexible datacenter. For example, remote master control **420** or critical datacenter control system **1022** within system **1000** may be configured to determine whether to route the computational operation to flexible datacenter **1002**. In other examples, a flexible datacenter control system **220** may determine whether to route the computational operation to flexible datacenter **1002**. For instance, the flexible datacenter control system **220** may determine whether the computing systems **100** have the availability to perform one or more computational operations within the queue system **1008**. In further examples, other components can perform the determination step.

Determining whether to route the computational operation to a flexible datacenter, such as flexible datacenter **200**, can involve considering various factors, such as a cost of execution to provide computing resources at the flexible datacenter relative to the cost of providing computing resources at the critical datacenter. The determination may also factor the availability of the flexible datacenter as well as the cost and availability of unutilized behind-the-meter power from one or more behind-the-meter sources. Other factors can be considered within examples, such as monitored power system conditions and operational directives.

At step **1206**, the method involves causing the computational operation to the flexible datacenter via a communication link, such as links **1006**, **425**, **1002**, **1010**, **1012a**, and/or **1012b**, based on a determination to route the computational operation to the flexible datacenter. Sending the computational operation may enable computing systems of the flexible datacenter to provide computing resources to fulfill the request.

In some examples, remote master control **420**, critical datacenter control system **1022**, or another component within system **1000** may determine that the identified computational operation should be routed to the critical datacenter **1004**. The determination may be based on various factors, such as a cost of execution to provide computing resources at the flexible datacenter relative to the cost of providing computing resources at the critical datacenter. The determination may also factor the availabilities of the critical datacenter **1004** and the flexible datacenter **200** as well as the

34

cost and availability of unutilized behind-the-meter power from one or more behind-the-meter sources. Other factors may be considered. As such, one or more components may route the computational operation to the critical datacenter **1004** to enable the computing systems **1020** to fulfill the computational request.

FIG. **13** illustrates a method for managing queue distribution between a critical datacenter and a plurality of flexible datacenter in accordance with one or more embodiments of the present invention. The method serves an example and may include other steps within other examples. The method of FIG. **13** is similar to the method of FIG. **12**, and steps, components, and aspects illustrated and/or described in FIG. **13** that are similar to or the same as components or aspects illustrated and/or described in FIG. **12** should be considered to have the same characteristics as previously illustrated and/or described.

At step **1302**, the method involves identifying, using a queue system, a computational operation to be performed. The computational operation may be performed at a critical datacenter, one or more flexible datacenters, or a combination of datacenters. In some examples, the queue system may be distributed into multiple queue subsystems with each queue subsystem organizing a set of computational operations according to rules associated with that queue subsystem. For example, the queue system may include a first queue subsystem and a second queue subsystem where each queue subsystem organizes computational operations accessibly potentially only by a subset of the plurality of flexible datacenters. For instance, a first set of flexible datacenters of the plurality of flexible may be configured to obtain computational operations from the first queue subsystem and a second set of flexible datacenters of the plurality of flexible datacenters are configured to obtain computational operations from the second queue subsystem.

At step **1304**, the method involves determining whether to route the computational operation to a flexible datacenter in a plurality of flexible datacenters. However, in particular, multiple flexible datacenters may be available to receive the computational operation. As such, a computing system, such as remote master control system **420** or critical datacenter control system **1022**, may determine whether to route the computational operation to a flexible datacenter out of the multiple available.

In some examples, the determination may be made by one or more datacenter control systems associated with the plurality of flexible datacenters. Each datacenter control system may determine whether or not its computing systems could currently handle the computational operation.

At step **1306**, the method involves, based on a determination to route the computational operation to a flexible datacenter in the plurality of flexible datacenters, determining a specific flexible datacenter in the plurality of flexible datacenters to route the computational operation to. The computing system may select a specific datacenter based on cost, availability, source of unutilized behind-the-meter power, or other factors. For example, the computing system may compare the cost associated with sending the computational operation to different flexible datacenters.

In some examples, a flexible datacenter or a critical datacenter may access and obtain the computational operation from the queue system. For example, a flexible datacenter from the plurality of flexible datacenters may obtain the computational operation upon determining that its computing systems are capable of supporting the computational operation (e.g., power is available, enough computing systems are free to operate on the computational operation).

35

At step 1308, the method involves causing the computational operation to be sent to the specific flexible datacenter via the communication link. Various components within the system may enable the computational operation to reach the specific flexible datacenter.

In further examples, the method described above may involve dividing the computational operation among multiple flexible datacenters.

Advantages of one or more embodiments of the present invention may include one or more of the following:

One or more embodiments of the present invention provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

One or more embodiments of the present invention allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive unutilized behind-the-meter power when it is available.

One or more embodiments of the present invention provide the use of a queue system to organize computational operations and enable efficient distribution of the computational operations to datacenters.

One or more embodiments of the present invention enable datacenters to access and obtain computational operations organized by a queue system.

One or more embodiments of the present invention allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

One or more embodiments of the present invention may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

One or more embodiments of the present invention may be powered by behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as distributed computing processes, with little to no energy cost.

One or more embodiments of the present invention provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit and/or generates incremental revenue.

One or more embodiments of the present invention allows for continued shunting of behind-the-meter power into a storage solution when a flexible datacenter cannot fully utilize excess generated behind-the-meter power.

One or more embodiments of the present invention allows for continued use of stored behind-the-meter power when a flexible datacenter can be operational but there is not an excess of generated behind-the-meter power.

It will also be recognized by the skilled worker that, in addition to improved efficiencies in controlling power delivery from intermittent generation sources, such as wind farms and solar panel arrays, to regulated power grids, the invention provides more economically efficient control and stability of such power grids in the implementation of the technical features as set forth herein.

While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that other embodiments may be devised that are within the scope

36

of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

What is claimed is:

1. A system comprising:

a flexible datacenter that is positioned to receive electrical power directly from a power generation system without the received power flowing through a grid comprising: a behind-the-meter (“BTM”) power input system configured to receive power from the power generation system prior to the power undergoing step-up transformation for transmission to the grid and provide the received power to the flexible datacenter instead of sending the received power to the grid;

a power distribution system; and

a plurality of computing systems powered by the BTM power input system via the power distribution system;

a queue system configured to organize a plurality of computational operations, wherein the queue system comprises multiple queue subsystems, and wherein each queue subsystem organizes a set of computational operations according to rules associated with the queue subsystem; and

a routing control system configured to:

determine, using the queue system, a computational operation to be performed;

determine an amount of BTM power that is available to the plurality of computing systems;

determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available; and

based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.

2. The system of claim 1, wherein the routing control system is a remote master control system positioned remotely from the flexible datacenter.

3. The system of claim 1, wherein the routing control system is a remote master control system positioned at the flexible datacenter.

4. The system of claim 1, wherein the flexible datacenter further comprises:

a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on one or more monitored power system conditions.

5. The system of claim 4, wherein the one or more monitored power system conditions comprises:

behind-the-meter (“BTM”) power availability at the plurality of computing systems.

6. The system of claim 4, wherein the one or more monitored power system conditions comprises:

one or more of excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation subject to economic curtailment, local power generation subject to reliability curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation.

7. The system of claim 1, wherein the flexible datacenter further comprises:

37

a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive.

8. The system of claim 7, wherein the operational directive originates from the routing control system and depends on behind-the-meter (“BTM”) power availability at the plurality of computing systems.

9. The system of claim 1, wherein the routing control system controls distribution of the plurality of computational operations from the queue system.

10. The system of claim 1, wherein a first queue subsystem of the queue system is located at the flexible datacenter.

11. The system of claim 10, further comprising:

a second datacenter, wherein a second queue subsystem of the queue system is located at the second datacenter.

12. The system of claim 11, wherein the second datacenter is a second flexible datacenter.

13. A flexible datacenter comprising:

a behind-the-meter (“BTM”) power input system configured to receive power from a power generation system prior to the power undergoing step-up transformation for transmission to a grid and provide the received power to the flexible datacenter instead of sending the received power to the grid;

a power distribution system;

a datacenter control system;

a plurality of computing systems, wherein the plurality of computing systems are powered by the BTM power input system via the power distribution system; and

38

a queue subsystem configured to organize a set of computational operations for performance by the plurality of computing systems, wherein the queue subsystem is part of a distributed queue system;

wherein the datacenter control system is configured to:

determine, using the queue system, a computational operation to be performed;

determine an amount of BTM power that is available to the plurality of computing systems;

determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available; and

based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.

14. The flexible datacenter of claim 13, wherein the distributed queue system includes at least two queue subsystems having different rules for organizing computational operations.

15. The flexible datacenter of claim 13, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive received from a remote master control system.

16. The flexible datacenter of claim 15, wherein the remote master control system is positioned remotely from the flexible datacenter.

* * * * *