

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSUNG ELECTRONICS CO., LTD.,
Petitioner

v.

NETLIST, INC.,
Patent Owner

Case IPR2025-01431
U.S. Patent No. 10,025,731

**DECLARATION OF STEVEN PRZYBYLSKI, PH.D. IN SUPPORT OF
PATENT OWNER'S RESPONSE**

Mail Stop "PATENT BOARD"

Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Netlist EX2058
Samsung v. Netlist
IPR2025-01431

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	MY BACKGROUND AND QUALIFICATIONS	2
III.	SUMMARY OF OPINIONS.....	4
IV.	LIST OF DOCUMENTS CONSIDERED	4
V.	LEGAL UNDERSTANDING.....	6
	A. My Understanding of Claim Construction.....	6
	B. My Understanding of Obviousness	7
	C. A Person of Ordinary Skill in the Art	9
VI.	THE '731 PATENT.....	10
	A. Overview	10
	B. Level of Skill in the Art.....	14
VII.	THE ASSERTED REFERENCES	15
	A. Ellsberry	15
	B. Dour.....	19
	C. Abadeer.....	23
VIII.	THE CLAIMS WOULD NOT HAVE BEEN OBVIOUS.....	24
	A. The Combinations of Grounds 1 and 2 do not render Claim 1 Obvious.....	24
	1. Ellsberry's memory controller cannot access an independent set of memory devices on the memory module (element 1.c).	24
	a. Memory controllers typically access ranks of a multi- rank DIMM using rank-select or chip-select signals.....	25
	b. Because Ellsberry's banks are "transparent" to the processing unit, the processing unit cannot access individual banks.....	28
	c. None of the previous Board decisions determined that Ellsberry discloses element 1.c.....	32
	2. The combination of Ellsberry and Dour does not render obvious dynamically controlling an impedance matching circuit based	

	on which rank of a multi-rank memory module is selected (element 1.f.2)).....	35
3.	A POSA would not have implemented Dour's ODT in Ellsberry's switch ASIC to terminate signals.	44
IX.	CONCLUSION.....	49

I, Steven Przybylski, Ph.D., hereby declare as follows.

I. INTRODUCTION

1. I have been retained as an expert witness on behalf of Patent Owner Netlist, Inc. (“Patent Owner”) for the above-captioned *inter partes* review (IPR). I am being compensated for my time in connection with this IPR at my standard consulting rate.

2. I understand that this Declaration accompanies the Patent Owner’s Response filed in this IPR involving U.S. Patent No. 10,025,731 (“the ’731 patent”) (EX1001), which resulted from U.S. Patent Application No. 14/715,491 (“the ’491 application”), filed on May 18, 2015. I understand that the ’731 patent claims a priority date of April 14, 2008. I refer to this date throughout this declaration.

3. In preparing this Declaration, I have reviewed the ’731 patent and each of the documents cited herein, in light of general knowledge in the art before April 14, 2008. In formulating my opinions, I have relied upon my experience, education, and knowledge in the relevant art. In formulating my opinions, I have also considered the viewpoint of a person of ordinary skill in the art (“POSA”) in the April 2008 timeframe. I understand that the ’731 patent claims priority to provisional patent applications that were filed in April 2008. I also understand that Samsung has not challenged this priority claim, and thus I did not perform any

analysis regarding the '731 patent's priority claims.

II. MY BACKGROUND AND QUALIFICATIONS

4. Below is a summary of my education and experience. My *curriculum vitae*, included as Exhibit 2059, records my education, experience, and publications in greater detail.

5. I earned a Bachelor of Applied Science from the University of Toronto in 1980. I was enrolled in the Engineering Science program, completing a course of study combining the Electrical Engineer and Computer Science options.

6. I earned a Master of Science in Electrical Engineering degree and a Ph.D. in Electrical Engineering in 1982 and 1988 respectively, both from Stanford University.

7. I also earned a Master of Business Administration from the Haas School of Business at the University of California at Berkeley in 2000.

8. I have extensive experience with memory semiconductor integrated circuits and the memory systems constructed of them. At Stanford, my dissertation was on the optimization of single- and multi-level cache hierarchies to maximize system-level performance. Also at Stanford, I was a member of the core team that architected, designed, built, and tested the seminal MIPS processor. I was responsible for the design of the instruction decode and control units as well as architecting the virtual memory support. I assembled and debugged the entire

microprocessor design and oversaw its fabrication and testing. In 1984 and 1985, I took a leave of absence from Stanford to become a member of the founding team of MIPS Computer Systems, a startup in California that designed, built, and sold processors and computer systems. In 1989, after finishing my doctorate and brief post-doctorate at Stanford, I returned to MIPS Computer Systems. Throughout my two periods of employment at MIPS Computer Systems, I had a number of titles and worked on a variety of projects. Noteworthy were stints as processor architect, systems architect, hardware-software team liaison, and Chief Scientist of the High-End Systems group. During the second period, I also served as a Consulting Assistant Professor in the Department of Electrical Engineering at Stanford. I designed and taught an advanced graduate level course on cache and memory system design.

9. In 1991, I left MIPS to become an independent consultant. In that capacity I have provided technical, strategic, marketing, and intellectual property services to a large number of clients both domestically and internationally.

10. I am a member of both the Institute of Electrical and Electronics Engineers (IEEE) and the Association of Computing Machinery (ACM). In addition, in 1997 I became a registered patent agent. I am listed as the inventor on 11 domestic patents and their foreign counterparts. The subject matters include memory devices, memory subsystems, networks, error detection and correction

codes, processors and cache memories. Between 1994 and approximately 2000, I was a leading independent analyst following, writing, and teaching about both the technical and business aspects of the DRAM industry. Overall, I have written three books and over 25 articles in the areas of memory, processors, and computer systems. In particular, I wrote and revised an 850-page book entitled “New DRAM Technologies: A Comprehensive Analysis of the New Architectures.” I have presented many lectures, classes and seminars on semiconductor memories and their memory systems, processors and computer systems. I have presented full and half-day seminars on discrete and embedded DRAMs to over 2,500 engineers, marketers, managers, and financial analysts.

III. SUMMARY OF OPINIONS

11. It is my opinion that claims 1-18 of the '731 patent would not have been obvious over the combinations that Samsung and Dr. Wolfe propose in Grounds 1 and 2.

IV. LIST OF DOCUMENTS CONSIDERED

Exhibit/Paper No.	Description
1	Petition for Post-Grant Review of U.S. Patent No. 10,025,731
1001	U.S. Patent No. 10,025,731

Exhibit/Paper No.	Description
1002	File History of U.S. Patent No. 10,025,731 (Application No. 14/715,491), <i>available at</i> "><https://patentcenter.uspto.gov/applications/14715491/ifw/docs?application=>
1003	Declaration of Dr. Andrew Wolfe
1004	<i>Curriculum Vitae</i> of Dr. Andrew Wolfe
1005	U.S. Patent App. Publication No. 2006/0277355 by Ellsberry et al.
1010	<u>JEDEC Standard 21-C</u> , DDR SDRAM Registered DIMM Design Specification (January 2002)
1018	U.S. Patent No. 7,020,818 to Dour et al.
1039	U.S. Patent Application Publication No. 2008/0218290 to <u>Abadeer</u> et al.
1062	Final Written Decision, <i>SK hynix Inc. et al. v. Netlist, Inc.</i> , IPR2018-00362 and -00363, Paper No. 29 (PTAB June 27, 2019) ('907 patent)
1066	Final Written Decision, <i>Samsung Electronics Co., Ltd. et al. v. Netlist, Inc.</i> , IPR2022-00639, Paper 45 (PTAB Oct. 18, 2023) ('339 patent)
1078	Final Written Decision, <i>Samsung Electronics Co., Ltd. et al. v. Netlist, Inc.</i> , IPR2022-00615, Paper 96 (PTAB Apr. 17, 2024) ('912 patent)
1093	Invalidity claim chart
2060	Transcript of the Deposition of Andrew Wolfe, Ph.D., May 15, 2026
2122	U.S. Patent No. 8,497,712 to Nguyen et al., issued July 30, 2013

Exhibit/Paper No.	Description
2123	Zhang, et al., “A Built-In-Test Circuit for Functional Verification & PVT Variations Monitoring of CMOS RF Circuits,” <i>ISSC 2006, Dublin Institute of Technology</i> , June 28-30, 2006, pp. 217-222
2124	Majzoub, et al., “PVT Variation Impact on Voltage Island Formation in MPSoC Design,” <i>10th Int’l Symposium on Quality Electronic Design</i> , March 16-18, 2009, pp. 814-819
2125	Gopal, et al., “Ultra Low-Supply Voltage Reference Generator with Low Sensitivity to PVT Variations,” <i>5th Asia Symposium on Quality Electronic Design</i> , August 26-28, 2013, pp. 18-21
2126	Miyashita, et al., “A Phase Noise Minimization of CMOS VCOs over Wide Tuning Range and Large PVT Variations,” <i>Proceedings of the IEEE 2005 Custom Integrated Circuits Conference, 2005</i> , September 21, 2005, pp. 583-586

V. LEGAL UNDERSTANDING

12. I have also relied upon various legal principles (as explained to me by counsel) in formulating my opinions. My understanding of these principles is summarized below.

A. My Understanding of Claim Construction

13. I understand that during an *inter partes* review proceeding, claims are to be construed in light of the specification as would be read by a person of ordinary skill in the relevant art at the time the application was filed. I understand that claim terms are given their ordinary and customary meaning as would be understood by a person of ordinary skill in the relevant art in the context of the

entire disclosure. A claim term, however, will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the claim term in the specification. In this case, the claim term will receive the definition set forth in the patent.

B. My Understanding of Obviousness

14. I understand that a patent claim is invalid if the claimed invention would have been obvious to a POSA at the time the application was filed. This means that even if all of the requirements of the claim cannot be found in a single prior art reference that would anticipate the claim, the claim can still be invalid.

15. To obtain a patent, a claimed invention must have, as of the priority date, been nonobvious in view of the prior art in the field. I understand that an invention is obvious when the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious to a POSA at the time the invention was made.

I understand that to prove that prior art or a combination of prior art renders a patent obvious, it is necessary to:

(1) identify the particular references that, singly or in combination, render the patent obvious;

(2) specifically identify which elements of the patent claim appear in each of the asserted references; and

(3) explain how the prior art references would have been combined or modified in order to create the inventions claimed in the asserted claim.

16. I also understand that prior art references can be combined under several different circumstances. For example, it is my understanding that one such circumstance is when a proposed combination of prior art references results in a system that represents a predictable variation, which is achieved using prior art elements according to their established functions.

17. I also understand that when considering the obviousness of a patent claim, one should consider whether a teaching, suggestion, or motivation to combine the references exists so as to avoid impermissibly applying hindsight when considering the prior art. I understand this test should not be rigidly applied, but that the test can be important to avoid such hindsight. I also understand that, when combining references, characterizing the references as “primary” and “secondary” is merely a matter of presentation with no legal significance. For instance, whether a combination is presented as reference A in view of reference B or reference B in view of reference A does not affect the underlying obviousness considerations for a patent claim.

18. I understand that certain objective indicia can be important evidence as to whether a patent is obvious or nonobvious. Such indicia include:
(1) commercial success of products covered by the patent claims; (2) a long-felt

need for the invention; (3) failed attempts by others to make the invention; (4) copying of the invention by others in the field; (5) unexpected results achieved by the invention as compared to the closest prior art; (6) praise of the invention by the infringer or others in the field; (7) the taking of licenses under the patent by others; (8) expressions of surprise by experts and those skilled in the art at the making of the invention; and (9) the patentee proceeded contrary to the accepted wisdom of the prior art.

19. I also understand that “obviousness” is a legal conclusion based on the underlying factual issues of the scope and content of the prior art, the differences between the claimed invention and the prior art, the level of ordinary skill in the prior art, and any objective indicia of non-obviousness.

20. For that reason, I am not rendering a legal opinion on the ultimate legal question of obviousness. Rather, my testimony addresses the underlying facts and factual analysis that would support a legal conclusion of obviousness or non-obviousness, and when I use the term obvious, I am referring to the perspective of one of ordinary skill in the art at the time of invention.

C. A Person of Ordinary Skill in the Art

21. I understand that a person of ordinary skill in the relevant art (“POSA”) is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity—not an automaton.

22. I have been asked to consider the level of ordinary skill in the field that someone would have had at the time the claimed invention was made. In deciding the level of ordinary skill, I considered the following:

- the levels of education and experience of persons working in the field;
- the types of problems encountered in the field; and
- the sophistication of the technology.

23. My opinion below explains how a POSA would have understood the technology described in the references I have identified herein around the April 14, 2008 timeframe, which I have been advised is the earliest possible effective filing date for the '731 patent.

24. Regardless of whether I use “I” or a “POSA” during my technical analysis below, all of my statements and opinions are to be understood to be based on how a POSA would have understood or read a document at the time of the alleged invention.

VI. THE '731 PATENT

A. Overview

25. The '731 patent describes a memory module that reduces signal noise and reflections caused by impedance mismatches between the transmission lines and circuits that receive the signals. *See* EX1001, 2:2-14, 4:39-43, 9:55-10:29.

Specifically, that '731 patent describes a memory module that includes memory devices arranged in multiple ranks and programmable impedance matching circuits that are dynamically controlled based on which of the multiple ranks is selected to communicate data with an external system memory controller. *See* EX1001, 5:65-6:28, 10:59-11:17.

26. Figure 3, reproduced below, illustrates an external memory controller 106 that communicates with a memory module 100 that includes memory devices 110. EX1001, 5:14-21. The memory module 100 may also include a circuit 112 that “comprises a correction circuit 150 (see, e.g., FIG. 4) which reduces noise in one or more signals” EX1001 5:22-40.




FIG. 3

EX1001, FIG. 3.

27. “Referring to FIG. 4, ... the correction circuit 150 includes one or more elements to provide for coarse correction of signal noise due to impedance matches (e.g., discontinuity regions) in the system. The at least one coarse correction element 152 comprises at least one programmable impedance matching circuit.” EX1001, 10:40-46, FIG. 4 (reproduced below).




FIG. 4

EX1001, FIG. 4.

28. The '731 patent explains that “impedance values of the coarse correction element 152 may be dynamically controllable, configurable, and/or programmable, so as to match the characteristic impedances of the transmission lines (e.g., signal paths) in the system. The values of these impedances may differ (e.g., may be programmed differently) based on which particular signal path is activated.” EX1001, 10:66-11:6. And, “[d]uring operation, one or more ranks of a memory module 100 may be selected or activated from the plurality of ranks by control signals that are received from a component of the system (e.g., a system memory controller 106 or a local memory controller of the memory module 100).

Examples of such control signals include, but are not limited to rank-select signals, also called chip-select signals.” EX1001, 6:20-28, FIG. 3 (reproduced above).

29. The ’731 patent thus describes a memory module that includes signal correction circuits, each including at least one programmable impedance matching circuit, to correct signals transmitted to and from the module’s memory devices, and control circuitry configured to dynamically control the at least one programmable impedance matching circuit. *See* EX1001, 9:9-13, 9:55-62. The claimed correction circuitry can reduce signal noise due to impedance mismatch in the system, enabling the claimed memory modules to provide stable, reliable signals at high data rates. *See* EX1001, 4:30-67. The claimed control circuitry can dynamically control impedance values of the programmable impedance matching circuit based on selection or activation of a rank of memory devices. *See* EX1001, 6:20-28; 10:66-11:6.

B. Level of Skill in the Art

30. A person of ordinary skill in the art of the ’731 patent (which I abbreviate “POSA”) would have had a Bachelor’s degree in Electrical Engineering, Computer Science, or Computer Engineering, or a related field, and at least two years of experience in the design, development, or testing of computer memory systems. Additional training can substitute for education or research experience, and vice versa.

VII. THE ASSERTED REFERENCES

31. Samsung presents two Grounds of unpatentability based on obviousness. In Ground 1, Samsung alleges that claims 1-18 are obvious over the combination of Ellsberry and Dour. In Ground 2, Samsung alleges that claims 1-18 are obvious over the combination of Ellsberry, Dour, and Abadeer. Below I provide a summary of the three references that Samsung used in Grounds 1 and 2.

A. Ellsberry

32. Ellsberry is directed to a capacity-expanding memory device. EX1005, Title. In the background, Ellsberry describes a conventional memory system. EX1005, ¶¶2-5. A memory module, such as a dual in-line memory module or “DIMM,” may include dynamic random-access memory (“DRAM”) chips that can store information. EX1005, ¶¶2-3. The memory module is inserted into “a slot” on a “computer motherboard,” enabling information to be exchanged between the memory module and a “system memory controller” via a memory bus on the motherboard. EX1005, ¶¶2-5.

33. In conventional systems, the system memory controller knows the configuration of the memory module(s) that it communicates with. *See* EX1005, ¶4. The system memory controller typically obtains the memory module’s configuration information by reading it from a serial presence detect (“SPD”) chip that is on the memory module. *See* EX1005, ¶34; EX2060, 32:17-33:11. And

because the system memory controller knows the memory module's configuration, the system memory controller can *access* memory locations on the memory module: "To *access* locations in memory devices, such as DRAM, a memory location is typically specified by the system over an address bus. This address is decoded to *access* the correct memory device." EX1005, ¶4 (emphasis added).

34. Ellsberry recognized that electronic devices were becoming "more sophisticated" and there was a "need for greater storage or memory increases." EX1005, ¶5. And because "many electronic applications are restricted by industry standards and physical limitations," it was "undesirable and impractical" to increase the size of the memory bus on the system's motherboard. EX1005, ¶5. Ellsberry thus sought to "expand[] the memory capacity of a memory module without increasing the bus size." EX1005, ¶26; *see also* EX1005, ¶10.

35. Ellsberry's solution was a memory module in which smaller-capacity memory devices "emulate" a single, higher capacity memory device, EX1005, ¶¶10, 26, while making the underlying architecture of the memory module "transparent" to the rest of the system, EX1005, ¶¶1, 27. For example, "two separate SDRAM DDR devices" can "appear as a single higher-capacity SDRAM DDR device to" the system memory controller while the actual configuration of the memory module is hidden to the system memory controller. EX1005, ¶26. Importantly, Ellsberry organizes the memory devices on the memory module into

“banks” but does not share this information with the system memory controller.

EX1005, ¶¶27, 30.

36. To hide the memory module’s architecture from the system memory controller, Ellsberry implemented additional logic on the memory module that “intercepts and modifies communications from the host computer.” EX1003, ¶51; *see also* EX2060, 42:21-43:5 (Dr. Wolfe explaining that “the system cannot tell logically or electrically that there is ranked multiplication going on, because there’s an intermediate circuit that isolates the multiple memory banks from the memory controller and performs all of the translations that -- that take place.”). Specifically, Ellsberry’s “capacity-expanding feature ... is accomplished by a combination of a control unit 204 and one or more memory bank switches 206 & 208,” which are shown in Figure 2 below. EX1005, ¶28.




Fig. 2

EX1005, FIG. 2.

37. Figure 2 depicts an embodiment of Ellsberry's capacity-expanding memory system. EX1005, ¶28. "The DIMM interface 202 may be coupled to a memory socket and communication bus over which data, memory addresses, commands, and control information are transmitted." EX1005, ¶28. Ellsberry explains that "control unit 204 receives memory addresses and commands over the DIMM interface 202" and the "memory bank switches 206 & 208 receive data information from the DIMM interface 202 via data buses 230 & 232, respectively." EX1005, ¶29. "Each memory bank switch 206 & 208 is communicatively coupled to a plurality of memory banks that may have one or more memory devices (e.g.,

dynamic random access memory (DRAM)).” EX1005, ¶30.

38. Ellsberry explains that the control unit 204 performs address space translation to identify which banks and memory locations to access. EX1005, ¶¶31, 35, 37, 39; FIG. 7; *see also* EX2060, 42:21-43:5. Because the memory module’s architecture is transparent, the system memory controller provides an address in a “primary space” that the control unit 204 maps or translates into a “secondary space.” EX1005, ¶¶35 (“The control unit 204 then handles mapping the logical memory addresses it receives via the DIMM interface 202 to a corresponding bank for a particular memory bank coupled to the switches 206 and 208.”), 37 (“[A] primary space address received by the control unit 204 is mapped to a secondary space address corresponding to the memory banks.”). In particular, “the control unit 204 decodes a memory address received over the DIMM interface 202, determines to which memory bank the received address corresponds, and causes the memory bank switch 206 and 208 to activate the correct memory bank.” EX1005, ¶31.

B. Dour

39. Dour is directed to programmable on-die termination (“ODT”) for properly terminating signals that travel on buses between components on a computer system’s main (or mother) board. Dour teaches that “[c]omputer system components include processors” and “memory modules” that “communicate with

one another over a set of buses.” EX1018, 1:15-18. “Buses are typically communication lines laid out on a printed circuit board such as the main board of a computer system.” EX1018, 1:30-32.

40. Figure 2 of Dour is reproduced below and illustrates communication line 205 (red) between a memory controller 213 and a memory module 211.

EX1018, 2:42-44, FIG. 2. “Communication line 205 may be part of a system bus such as a front side bus, a memory bus, or similar communication line.” EX1018, 2:54-56.




FIG. 2

EX1018, FIG. 2 (annotated).

41. Dour explains that “memory module 211 may be a dynamic random access memory [(DRAM)] module” that is “packaged as a dual in line memory

module (DIMM).” EX1018, 2:46-49. Dour further explains that “[m]emory module 211 may be a double data rate (DDR) module” EX1018, 2:50-51. But other than describing the memory module as a DDR DRAM packaged as a DIMM, Dour provides no further details of the memory module. Relevant here, Dour does not disclose that the memory module is a multi-rank memory module. Nor does Dour suggest on die termination is used to terminate signals that are transmitted within the memory module itself.

42. Dour teaches programming the ODT based on certain conditions *in the device* that receives a signal (such as a memory controller), not based on the length of the communication line. EX1018, 3:60-65, 4:4-7, 4:7-11, 4:11-15, 4:17-20, 4:21-22, 5:20-23, 6:50-53. Specifically, Figure 4 of Dour is reproduced below and shows a memory controller that includes a “testing circuit 401 to determine conditions *in the device*.” EX1018, 3:60-61 (emphasis added).




FIG. 4

EX1018, FIG. 4 (annotated).

43. Testing circuit 401 includes a circuit block 403, a finite state machine (“FSM”) 405, and a feedback loop 409. EX1018, 3:65-4:3. Dour teaches that the FSM 405 encodes the “conditions of circuit block 403” into a signal that is used to adjust the ODT. EX1018, 4:4-20 (emphasis added). Dour explains that these “conditions ... include the effects of the physical characteristics of the device, such as quality and type of manufacture, materials and similar physical characteristics, temperature, voltage, and similar device conditions have on [sic] the signal processed by circuit block 403.” EX1018, 4:7-12 (emphasis added); *see also*

EX1018, 4:12-15 (“In one embodiment, finite state machine 405 may encode the impedance of signals in block 403 based on the physical characteristics, voltage levels, and temperatures (PVT) *of block 403.*”), 6:50-53 (“[A] testing circuit or similar component detects the physical conditions, voltage conditions and temperature conditions (PVT) *of the component or on die termination circuit* (block 701).”), FIG. 7 (emphasis added). Nowhere does Dour teach or suggest that the ODT is adjusted based on the characteristics of the communication line or bus.

C. Abadeer

44. Abadeer is titled “Automatic Driver/Transmission Line/Receiver Impedance Matching Circuitry.” EX1039, Title. It is principally about the challenges related to operating CMOS circuits at very high frequencies. In the 2007 timeframe, gigahertz frequencies were challenging to attain. In particular, for data communication, the mismatch between reactive component of the impedance between the transmitter and receiver produced voltage transients that degraded reliability. The solution proposed was to match impedances automatically as a function of operating frequency. The principal feature of the solution was to match the operating frequency of the data communication with a PLL. The VCO of the PLL drives varactors (i.e. variable capacitance capacitors) the matching capacitor tracks with frequency.

45. Abadeer is cited in the Argument portion of the Petition with regards

to claims and elements 1.e.2, 4, and 16, which I do not address directly below.

VIII. THE CLAIMS WOULD NOT HAVE BEEN OBVIOUS.

46. The claims of '731 patent are challenged on two grounds in the Petition. In Ground 1, Claims 1-18 are challenged based on the combination of Ellsberry and Dour, rendering each of them obvious.

47. In Ground 2, Claims 1-18 are additionally challenged based on the combination of Ellsberry, Dour, and Abadeer, rendering each of them obvious.

A. The Combinations of Grounds 1 and 2 do not render Claim 1 Obvious

1. Ellsberry's memory controller cannot access an independent set of memory devices on the memory module (element 1.c).

48. Claim element 1.c requires a “plurality of memory devices” that are “arranged in multiple ranks” in which “each rank compris[es] an independent set of memory devices *that can be accessed by the system memory controller* to access a full bit-width of the memory bus.” (emphasis added). The plain and ordinary meaning of this claim language is that the memory controller must be able to access each rank (i.e., each independent set of memory devices that form a rank) of a multi-rank memory module.

49. Samsung relies on Ellsberry for element 1.c. Pet., 62-64; EX1003, ¶¶128-130. Neither Dour nor Abadeer disclose a memory module that includes ranks. EX2060, 74:16-74:7 (admitting that Dour does not disclose ranks), 21:6-16

(admitting that Abadeer does not disclose ranks). So Ellsberry is the only reference in Grounds 1 and 2 that is relevant to element 1.c. *See* EX1003, ¶¶128-130 (relying only on Ellsberry to address element 1.c); EX2060, 22:4-10.

50. But Ellsberry intentionally hides the bank arrangement and architecture of the memory module from the host system (including the system memory controller). Namely, Samsung alleges that Ellsberry’s “processing unit 102” is the claimed “memory controller” and that Ellsberry’s physical banks on the memory module are the claimed “ranks.” Pet., 58, 62-64. But Ellsberry teaches that the details of the memory module—particularly, the arrangement of its physical banks—are transparent to the processing unit (or system memory controller). Because Ellsberry’s banks are intentionally hidden from the rest of the system (including the processing unit, operating system, memory controller, etc.), Ellsberry’s system memory controller cannot access each of the individual banks of the memory module, as required.

a. Memory controllers typically access ranks of a multi-rank DIMM using rank-select or chip-select signals.

51. The ’731 patent explains that the memory controller can access a specific rank on the multi-rank memory module using specific control signals, such as rank-select or chip-select signals. Specifically, “[d]uring operation, one or more ranks of a memory module 100 may be selected or activated from the plurality of ranks by control signals that are received from a component of the system (e.g., a

system memory controller 106 or a local memory controller of the memory module 100).” The ’731 patent also explains that “such control signals include, but are not limited to, rank-select signals, also called chip-select signals.” EX1001, 6:20-26.

52. The ’731 patent’s description of how memory controllers can access a rank on a multi-rank DIMM is consistent with contemporary industry standards. EX1003, ¶51 (explaining that “a ‘rank’ (sometimes called a ‘bank’) is an independent set of one or more memory devices on a memory module that act together in response to command signals, including *chip select signals*, to read or write the full bit-width of the memory module.”) (emphasis added). For example, the JEDEC specification for DDR SDRAM Registered DIMM Design provides the pinout of standard memory module (DIMM) during the relevant timeframe. EX1010, 6-9; *see also* EX2060, 27:8-33:21. The DIMM includes pins, shown below, that are described as “chip select lines” that receive commands sent by a system memory controller to select physical banks on the DIMM. EX1010, 6; EX2060, 30:9-31:4, 32:2-16.

$\overline{S}0 - \overline{S}3$	SDRAM chip select lines (Physical. banks 0, 1, 2, and 3)
---------------------------------	---

EX1010, 6.

53. In typical systems, the memory controller reads configuration information from a Serial Presence Detect (SPD) chip on the DIMM, which may

include the number of physical banks on the DIMM. EX2060, 32:17-33:11;
EX1010, 68. The JEDEC specification for DDR SDRAM Registered DIMM
Design provides the example below, which shows an SPD that includes 256 bytes
of configuration information. EX1010, 68. Byte 5 indicates that the “Number of
Physical Banks” on this example DIMM is “2.” *Id.* The memory controller can use
this information to access specific banks on the DIMM via chip select signals.
EX2060, 32:17-33:11.

Serial Presence Detect Example Raw Card Version ‘A’ (Part 1 of 3)

64M x 72 DDR

Byte #	Description	SPD Entry Value	Serial PD Data Entry (Hexadecimal)	Notes
0	Number of Serial PD Bytes Written during Production	128	80	
1	Total number of bytes in Serial PD Device	256	08	
2	Fundamental Memory Type	SDRAM DDR	07	
3	Number of Row Addresses on Assembly	13	0D	
4	Number of Column Addresses on Assembly	10	0A	
5	Number of Physical Banks on DIMM	2	02	

EX1010, 68 (reproduced in-part and annotated).

54. By contrast, Ellsberry intentionally hides the memory module’s bank-related information from the memory controller, emphasizing that the capacity-expanding aspects of its memory module are “transparent” to the rest of the system. EX1005, ¶¶1, 27. Indeed, Ellsberry discloses that the memory module may include an SPD but discloses only that the SPD contains information about the capacity of the memory module, not details of the bank arrangement. EX1005, ¶34. Specifically, Ellsberry teaches that the SPD contains information about “the total

module memory capacity and data width along with the individual emulated memory device capacity that is formed by two or more small-capacity memory devices.” *Id.* In other words, Ellsberry’s SPD does not report the actual memory devices but the characteristics of the emulated memory devices in aggregate. EX1005, ¶34. Indeed, the Ellsberry module includes an SPD that, in concept, reports the characteristics of a standard module the Ellsberry module emulates with half the actual number of memory devices on the module. But nowhere does Ellsberry state that information about the banks is included on the SPD. *Id.* And given Ellsberry’s emphasis on transparency and fidelity to the existing module standards, a POSA would have understood that Ellsberry’s SPD does not contain information about the memory banks.

b. Because Ellsberry’s banks are “transparent” to the processing unit, the processing unit cannot access individual banks.

55. Ellsberry does not disclose a memory controller that can access each rank (i.e., each independent set of memory devices that form a rank) in a multi-rank memory module. Instead, Ellsberry discloses that such details of the memory module are “transparent” (i.e., invisible) to the memory controller. EX1005, ¶¶1 (“a controller and architecture that permits transparent bank switching of memory devices”), 27 (“the capacity-expanding memory device 108 is compatible with existing system architectures and transparent to the rest of the system (e.g.,

microprocessor, operating system, etc.)”). Ellsberry’s control unit 204 of Figure 2 (or controller 510 of Figure 5), not the external processing unit/memory controller, accesses the specific banks on the memory module following address translation.

56. Samsung and Dr. Wolfe acknowledge that Ellsberry teaches “transparent” bank switching of memory devices. Pet., 29; EX1003, ¶¶92. They allege that Ellsberry’s disclosure that “the operating system addresses a single logical memory bank which is mapped by the controller 510 to the physical banks of memory devices” teaches a memory controller that can access individual ranks. Pet., 62; EX1003, ¶¶128. That is wrong; Ellsberry’s teaches the opposite.

57. As explained above, the crux of Ellsberry is that the “rest of the system” (host system, operating system, system memory controller, etc.) does not know that the memory devices on the memory module are half the reported size and that they are organized into banks. EX1005, ¶¶1, 10, 26-27. Those details are left to and handled by the logic on the memory module (e.g., the control unit 204 and memory bank switch 206 & 208 of Figure 2). EX1005, ¶¶28-31. And because these details are transparent to Ellsberry’s host system, it cannot access an “independent set of memory devices” that form an individual, physical bank. Instead, the host system sends an address in a primary address space (corresponding to the capacity of the memory module) that the logic on the memory module translates into a secondary address space and determines which

bank the logic should access. EX1005, ¶¶31, 37, 39-40, FIG. 7. At bottom, Ellsberry's memory module presents to the memory controller the characteristics of a standardized memory module logically, physically, electrically, and temporally. Thus, it is impossible for the memory controller to tell that it is not talking to a standardized memory module.

58. On cross-examination, Dr. Wolfe confirmed this operation in Ellsberry. Specifically, he confirmed that (1) Ellsberry's external memory controller "cannot tell logically or electronically" that the memory module includes multiple banks and (2) Ellsberry's control unit 204 ("an intermediate circuit") performs address translations in order to access the memory banks:

Q. What does Ellsberry mean that the capacity-expanding memory device is transparent to the rest of the system?


A. It means that *the system cannot tell logically or electrically* that there is ranked multiplication going on, because *there's an intermediate circuit that isolates the multiple memory banks from the memory controller and performs all of the translations* that -- that take place.

EX2060, 42:11-43:5 (emphasis added).

59. Dr. Wolfe's annotations to Ellsberry's Figure 2 (reproduced below) and his corresponding discussion of how Ellsberry's system operates further confirm that Ellsberry's system memory controller cannot access each of the independent banks. EX1003, ¶139; EX2060, 98:13-99:8 (discussing figure on page

101 of his declaration). As shown, Ellsberry's control unit 204 (labeled "Control ASIC") receives address and command signals from the DIMM interface 202.

EX1003, ¶139. The received address is an address in a primary address space that corresponds to the overall capacity of the memory module and thus corresponds to all of memory banks. EX1005, ¶¶37, 39; FIG. 7A (2nd column). The Control ASIC uses the primary address "to determine which rank ("Bank") to target for Read or Write command[s] via bus 220," then generates and sends chip select signals ("CS" highlighted in blue) via the bus 220 (annotated in blue) to access an independent bank. EX1003, ¶139 (quote in annotated figure); EX2060, 98:20-22 ("...so the command that comes into the control ASIC is not necessarily the command that gets sent out. It generates new commands.").



EX1003, ¶139 (p. 101)

60. Ellsberry's system memory controller thus sends address signals using a primary address space that corresponds to the overall capacity of the memory module defined by the sum of the capacity of the individual banks. The system memory controller cannot access the banks individually. Only logic on Ellsberry's memory module can access banks individually via chip select signals that are generated by the logic on the memory module.

c. None of the previous Board decisions determined that Ellsberry discloses element 1.c.

61. Samsung and Dr. Wolfe are incorrect that the Board has already

determined that Ellsberry discloses Element 1.c. Pet., 23, 64; EX1003, ¶¶86, 129.

Each of the previous final written decisions they rely on evaluated a patent from a different family than the '731 patent that had a different disclosure and different claims. Element 1.c is materially different from other claim elements that the Board evaluated, and thus the previous decisions do not control.

62. The final written decisions in IPR2018-00362 and IPR2018-00363 evaluated U.S. 9,606,907 (“the '907 patent”). EX1062. The '907 patent is not in the '731 patent family, *see* EX1001, (63), and does not recite the term “rank” nor that “each rank compris[es] an independent set of memory devices that can be accessed by the system memory controller to access a full bit-width of the memory bus.” *See* EX1062, 17, 19-21.

'731 patent, element 1.c	'907 patent, cited elements of claim 1
“a plurality of memory devices on the printed circuit board, the plurality of memory devices being arranged in multiple ranks, each rank comprising an independent set of memory devices that can be accessed by the system memory controller to access a full bit-width of the memory bus”	“A memory module having a width of N bits and configured to communicate with a memory controller via a set of control signal lines and M sets of n data lines, where M is greater than one and $N=M \times n$ ” “a plurality of memory devices coupled to the module control circuit, the plurality of memory devices including

	first memory devices and second memory devices, wherein, in response to the first module control signals, the first memory devices output or receive each N-bit wide data signal associated with the memory read or write command while the second memory devices do not output or receive any data associated with the memory read or write command”
--	---

63. The final written decision in IPR2022-00639 evaluated U.S. 10,949,339 (“the ’339 patent”). EX1066. The ’339 patent is not in the ’731 patent family, *see* EX1001, (63), and does not recite that “each rank compris[es] an independent set of memory devices that can be accessed by the system memory controller to access a full bit-width of the memory bus.” *See* EX1066, 33.

’731 patent, element 1.c	’339 patent, cited elements of claim 1
“a plurality of memory devices on the printed circuit board, the plurality of memory devices being arranged in multiple ranks, each rank comprising an independent set of memory devices that can be accessed by the system memory controller to access a full bit-	“double data rate dynamic random access memory (DDR DRAM) devices coupled to the PCB and arranged in multiple N-bit-wide ranks”

width of the memory bus”	
--------------------------	--

64. The final written decision in IPR2022-00615 evaluated U.S. 7,619,912 (“the ’912 patent). EX1078. The ’339 patent is not in the ’731 patent family, *see* EX1001, (63), and does not recite that “each rank compris[es] an independent set of memory devices that can be accessed by the system memory controller to access a full bit-width of the memory bus.” *See* EX1078, 38-39.

’731 patent, element 1.c	’912 patent, cited elements of claim 16
“a plurality of memory devices on the printed circuit board, the plurality of memory devices being arranged in multiple ranks, each rank comprising an independent set of memory devices that can be accessed by the system memory controller to access a full bit-width of the memory bus”	“a plurality of double-data-rate (DDR) memory devices coupled to the printed circuit board, the plurality of DDR memory devices having a first number of DDR memory devices arranged in a first number of ranks”

2. The combination of Ellsberry and Dour does not render obvious dynamically controlling an impedance matching circuit based on which rank of a multi-rank memory module is selected (element 1.f.2)).

65. Element 1.f.2 recites “wherein *the control circuitry is further configured to dynamically control the at least one programmable impedance matching circuit* in the each correction circuit of the set of correction circuits in the

each circuit of the at least one circuit *based on which of the multiple ranks is selected to communicate data with the system memory controller* in response to the address and control signals.”

66. Samsung does not allege that any of the references in Grounds 1 and 2 teach dynamically controlling an impedance matching circuit on a multi-rank memory module based on which rank is selected. Pet., 74-77; *see also* EX1003, ¶¶137-140; EX2060, 20:8-14, 76:4-15. Indeed, Dr. Wolfe confirmed that neither Ellsberry, nor Dour, nor Abadeer discloses this feature. EX2060, 76:16-77:6 (alleging that Ellsberry suggests impedance calibration “because of variation in wiring characteristics in the DIMM”), 76:7-10 (confirming that Dour does not discuss ranks), 24:14-17 (confirming that Abadeer does not discuss ranks). Accordingly, there does not appear to be any dispute that none of the references discloses a key element of the claimed invention: the ability to dynamically control impedance based on the specific rank that is selected.

67. Instead, Samsung alleges the combination of Ellsberry and Dour renders obvious element 1.f.2. Pet., 74-77; EX1003, ¶¶137-140; EX2060, 76:16-77:6. Specifically, Samsung posits that (1) it would have been obvious to implement “Dour’s ‘programmable on die termination circuit’ ... in Ellsberry’s ‘bidirectional signal drivers 402 & 404,’” Pet., 74, and (2) this combination would be further modified to dynamically calibrate impedance because of “variation of

the surrounding conditions of incoming data signals” depending on which bank is selected. Pet., 76-77.

68. Samsung’s basis for the latter modification is that Dour allegedly discloses dynamic impedance control based on the length of the communication line that carries the signal to be terminated, which Samsung alleges corresponds to which of Ellsberry’s banks is selected. Pet., 75-76 (alleging that “data lines to [Ellsberry’s] Bank 0/1 may be shorter than data lines to Bank 2/3, which results in smaller loads for the shorter data lines”); EX2060, 76:16-77:20 (unable to find any discussion of “wire lengths” in Dour and relying instead on “physical characteristics of the device”), 79:4-83:16 (discussing “PVT” in Dour). But Samsung misinterprets and mischaracterizes Dour’s teachings, and thus the modification fails.

69. Dour does not teach dynamic impedance control (or programmable ODT) based on wire length. *See* EX2060, 77:7-20. Indeed, it is a well-recognized fact that when terminating transmission lines, the length of the line is irrelevant to the appropriate termination value and that only the characteristic impedance of the line matters. EX1003, ¶¶52-54. Instead of line length, Dour teaches adjusting the ODT based on certain conditions *in the device* that receives and terminates the incoming signal (e.g., a memory controller). EX1018, 3:60-65, 4:4-7, 4:7-11, 4:11-15, 4:17-20, 4:21-22, 5:20-23, 6:50-53; EX2060, 78:3-10, 80:4-11, 81:1-4.

Specifically, Figure 4 of Dour is reproduced below and shows a memory controller that includes a “testing circuit 401 to determine conditions *in the device*.” EX1018, 3:60-61 (emphasis added). Testing circuit 401 includes a circuit block 403, a finite state machine (“FSM”) 405, and a feedback loop 409. EX1018, 3:65-4:3.




FIG. 4

EX1018, FIG. 4 (annotated).

70. Dour teaches that the FSM 405 encodes the “conditions of *circuit block 430*” into a signal that is used to adjust the ODT. EX1018, 4:4-20 (emphasis added); *see* EX2060, 79:7-81:4. Dour explains that these “conditions ... include the effects of the physical characteristics of *the device*, such as quality and type of

manufacture, materials and similar physical characteristics, temperature, voltage, and similar *device conditions* have on [sic] the signal processed by block 403.” EX1018, 4:7-12 (emphasis added); *see also* EX1018, 4:12-15 (“In one embodiment, finite state machine 405 may encode the impedance of signals in block 403 based on physical characteristics, voltage levels, and temperatures (PVT) *of block 403.*”) (emphasis added), 6:50-53 (“a testing circuit or similar component detects the physical conditions, voltage conditions and temperature conditions (PVT) *of the component or on die termination circuit* (block 701).”) (emphasis added), FIG. 7.

71. Nowhere does Dour teach or suggest that the ODT is adjusted based on the length of the communication line that carries the signal to be terminated or other characteristics of the communication line. Dr. Wolfe confirmed during cross-examination that Dour does not disclose programmable ODT based on “wire length.” EX2060, 77:7-20. Rather, he speculated that Dour’s disclosure of “environmental conditions” and “physical characteristics” include characteristics of the communication line 205 (i.e., the communication line on the motherboard between the memory controller and the memory module). EX2060, 78:11-16, 79:19-80:3 (“...I think it includes the characteristics of the communication line 205.”). But when confronted with the actual teachings of Dour related to PVT, he confirmed that Dour is referring to “the device” that terminates the signal.

EX2060, 78:3-10, 80:4-11, 81:1-4.

72. The concept of PVT had an established meaning in the industry and a POSA would have understood Dour's teachings to be consistent with this meaning. Specifically, PVT is used to refer to the process, (supply) voltage, and temperature of a device or circuit. Process refers to the process used to fabricate the device, which dictates the device's physical characteristics. Voltage and temperature refer to the supply voltage and the operating temperature of the device, which are environmental or operating conditions of the device. Accordingly, Dour's teachings regarding PVT relate to the terminating device not the communication line.

73. The established meaning of PVT to mean Process, (supply) Voltage, and Temperature is confirmed by many references. Consider U.S. 8,497,712 ("Techniques For Measuring Voltages in a Circuit", Filed 4/17/2008) which has a lot in common with Dour. EX2122, 3:1-10, 4:36-49, 5:31-38 ("The N reference voltages generated from resistor R3 and the N reference voltages generated from resistor R4 also remain at substantially the same voltages over an expected range of process, supply voltage, and temperature (PVT) variations of the integrated circuit..."), 8:17-26 ("An offset voltage in comparator 202 can be caused by process and layout variations in the integrated circuit die."), 11:63-12:4, 12:20-23 ("The band gap voltage VBG generated by the band gap reference voltage

generator remains substantially constant over an expected range of process, supply voltage, and temperature variations of the integrated circuit.”), 12:23-31, 16:40-45. This patent describes using measurement circuits, comparators, and an FSM to generate to generate calibration signals that compensate the performance of other circuits for variations in process, supply voltage, and temperature. EX2122, 7:34-52.

74. Across the literature of analog and digital systems, the term PVT shows up repeatedly always meaning, process, (supply) voltage, and temperature—the principal fabrication and environmental factors that influence circuit performance. A few of these myriad papers are presented here:

- EX2123, “A Built-In-Test Circuit for Functional Verification & PVT Variations Monitoring of CMOS RF Circuits”, Title, Abstract, p. 218
- EX2124, “PVT Variation Impact on Voltage Island Formation in MPSoC Design”, Title, Abstract, Introduction, p. 814.
- EX2125, “Ultra Low-Supply Voltage Reference Generator with Low Sensitivity to PVT Variations”, Title, Abstract, Introduction.
- EX2126, “A Phase Noise Minimization of CMOS VCOs over Wide Tuning Range and Large PVT Variations”, Title, Abstract, Introduction.

75. Thus, a POSA reading Dour’s description and use of “PVT” would understand Dour to be describing PVT as the characteristics and environmental

variables that influence on-die circuit performance, as is conventional. The correspondence between physical characteristics and process variation is strong. It is true that many authors at least sometime leave off “supply” from “supply voltage” when describing PVT problems and solutions, but that is only because they generally recognize that “supply” is understood and the use of voltage instead of supply voltage is simply a matter of linguistics efficiency. The normal scope of PVT as understood by a POSA thus clear: on-chip process variation, on-chip supply voltage variation, and on-chip temperature. This is entirely consistent with Dour’s statements about PVT. Specifically, the Petition cites Dour as teaching that “[t]hese conditions may change dynamically and driver signal 501 may dynamically alter the resistance of the on die termination circuit 500 accordingly.” Pet. 35 (citing EX1018, 5:23-25). But this fragment is part of a larger passage:

“Driver signal 501 may be based on the conditions of the device and thereby adjust the resistance to properly terminate an incoming signal based on the conditions of the device. These conditions may change dynamically and driver signal 501 may dynamically alter the resistance of the on die termination circuit 500 accordingly to allow transistor sets 503, 505 to act as a compensation circuit.”

EX1018, 5:20-27.

76. Thus, it is clear that the “conditions” that Dour is talking about are

on-chip conditions and that the driver signal 501 may cause the incoming signal to be properly terminated based on those conditions of the device. Thus, when on pages 54 and 53 of the Petition, the combination of Ellsberry and Dour is motivated, Petitioner oversteps. Petitioner does not suggest that Ellsberry teaches putting termination on the memory device side of the Switch ASIC. Pet., 50. Despite Petitioner's suggestion that Dour's "on die termination circuit may be programmed to compensate for environmental conditions" where the presumed environmental conditions are outside the "conditions in the device," there is no support for this idea within Dour or Ellsberry. Pet., 50, 52. Dour in particular teaches adjusting the termination circuit to account for on-device PVT changes, not changes in the characteristic impedance of the transmission line or the driver at the far end.

77. Petitioner does not suggest that such an extension of the teaching of Dour would have been obvious to a POSA. They only suggest that a POSA would have been motivated to implement Dour's on-die termination as disclosed in Ellsberry's Switch ASICs. Pet., 50. But that alone is insufficient because it does not disclose the final portion of element 1.f.2: "... dynamically control the at least one programmable impedance matching circuit in the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit based on which of the multiple ranks is selected."

78. Moreover, even if Dr. Wolfe was correct (which he is not) that the “physical characteristics” that Dour mentions is of the communication line 205, that is the communication line between the external memory controller and the memory module, as shown in Figure 2 of Dour below. EX1018, FIG. 2. Applied to Ellsberry, Dour’s teaching corresponds to the communication line between the system memory controller and Ellsberry’s control ASIC (for address and command signals) or Ellsberry’s switch ASIC (for data signals), not the communication line between the switch ASIC and the memory devices/banks.




FIG. 2

EX1018, FIG. 2 (annotated).

3. A POSA would not have implemented Dour’s ODT in Ellsberry’s switch ASIC to terminate signals.

79. Samsung argues that increases in the operating speed of memory

devices would have motivated a POSA to implement Dour's ODT circuitry in Ellsberry's switch ASICs. *See, e.g.*, Pet., 47 (“[A] POSITA would have understood the importance of proper signal termination as the speed of memory operations increases.”), 49 (“[I]n order to provide reliable data transmission at the higher transmission speeds expected in 2008.”), 52 (“To achieve such ‘higher communication speeds,’ a POSITA would have been motivated to use Dour's ODT functionality in Ellsberry's Switch ASICs.”), 75 (addressing element 1.f.2 and arguing that switching between Ellsberry's banks “requires accurate, dynamic termination, especially at a high operating speed”). Specifically, Samsung proposes implementing (1) Dour's programmable ODT circuit in Ellsberry's bidirectional signal drivers 402 & 404 *and* (2) Dour's testing circuit 402 in Ellsberry's Control Block of Figure 4 (which Samsung colors gray). Pet., 71-72.

80. I disagree that a POSA would have implemented Dour's PVT circuitry in Ellsberry's switch ASICs because Dour's Test circuit, finite state machine and ODT circuitry would have been too slow to handle high operating speeds as suggested by the Petitioner's expert. EX1003, ¶139.

81. Dour's testing circuit 401 includes a circuit block 403, a finite state machine (“FSM”) 405, and a feedback loop 409. EX1018, 3:65-4:3. Dour teaches that the FSM 405 encodes the “conditions of circuit block 403” into a signal that is used to adjust the ODT. EX1018, 4:4-20. These “conditions ... include the effects

of the physical characteristics of the device, such as quality and type of manufacture, materials and similar physical characteristics, temperature, voltage, and similar device conditions have on *[sic]* the signal processed by block 403.” EX1018, 4:7-12; *see also* EX1018, 4:12-15 (“In one embodiment, finite state machine 405 may encode the impedance of signals in block 403 based on the physical characteristics, voltage levels, and temperatures (PVT) of block 403.”), 6:50-53 (“[A] testing circuit or similar component detects the physical conditions, voltage conditions and temperature conditions (PVT) of the component or on die termination circuit (block 701).”), FIG. 7.

82. A POSA would have understood that the PVT characteristics of a device typically change slowly. Accordingly, a POSA would have known that Dour’s teaching of tracking and encoding PVT conditions into a driver signal that controls the programmable impedance circuit 500 does not have fast response. To start, the feedback loop 409 that iterates toward a final value includes the finite state machine 405, which is a clocked circuit that nominally takes at least one clock cycle to determine its output. The feedback loop 409 also includes the test block and two chip boundary crossings, which would have added additional latency. EX1018, FIG. 5. Further, Figure 6 of Dour shows the flowchart for the disclosed process for programming an on-die termination circuit. This flow chart includes a multiplier 601, which is an intrinsically complex operation. EX1018,

5:45-63. And Figure 7 of Dour illustrates the disclosed process for programming an ODT circuit to adjust for conditions in a device. EX1018, 6:48-56. Figure 7 confirms that Dour's ODT circuit may take multiple clock cycles to reach step 707 with a new stable termination value.

83. The bottom line is that Dour's ODT circuitry was designed to identify and react to PVT conditions that change much slower than the high-speed bank switching required by Ellsberry. Knowing this, a POSA would not have been motivated to make the combination that Samsung proposed. In fact, the significant delay that would be introduced by Dour's ODT circuitry would have dissuaded a POSA from making this combination.

84. Ellsberry's lines are also too short to justify adding Dour's ODT circuitry to Ellsberry's switch ASICs, as Samsung proposes. Specifically, a POSA would not have been motivated to include Dour's ODT circuitry in Ellsberry's switch ASICs to terminate signals transmitted between Ellsberry's switch ASICs and memory devices because the length of the wires between the switch ASICs and memory devices is too short to be negatively affected by impedance mismatch. The '731 patent explains that signal wires should be considered "transmission lines" that require proper termination if the length of the wire is greater than or equal to a fraction of the wavelength of the signal, such as $1/2$, $1/8$, or $1/10$ of the wavelength. EX1001, 9:63-10:6. This teaching was a well-known "rule of thumb"

in the industry that a POSA would have known about and followed.

85. Ellsberry discloses that “the control unit and bank switch combination may support SDRAM DDR1 and DDR2 specifications (incorporated herein by reference) with memory capacities of 256 Mb, 512 Mb, and 1 Gb in speed grades of 200/266/333/400 Mbps in DDR1 and capacities of 256 Mb, 512 Mb, 1 Gb, and 2 Gb in speed grades of 400/533 Mbps in DDR2.” EX1005, ¶46. Accordingly, the highest operating speed that Ellsberry contemplates is 533 Mbps (Mbps means megabits per second per pin). Because the DRAMs are double data rate (DDR), the clock frequency is half 533 Mbps, or 266.5 Mhz. The wavelength of a signal on the surface of an FR4 PCB is dependent on the speed of that signal and its frequency. The dielectric constant of FR4 is about 4.3, which makes the velocity of these signals to be about 0.5 c or 1.45×10^8 m/s. More conveniently, this is equivalent to about 170 ps/inch. The wavelength is equal to v/f . So, the wavelength of a 266.5 Mhz signal is about .55 m or 550 mm. And 1/10 of this wavelength (the smallest that the '731 suggests would require signal termination) is 55 mm.

86. Dr. Wolfe testified that the length of the longest wire between Ellsberry's switch ASIC and DRAM chips would be about 30 mm. EX2060, 58:18-59:8. Specifically, he indicated that the wire shown in red below would be about 15 mm, and the wire shown in blue below would be about 30 mm. *Id.* Because 30 mm is less than 55 mm (or 1/10 the wavelength of a signal at

Ellsberry's fastest operating speed), a POSA would have understood that no signal termination is required between Ellsberry's switch ASICs and memory devices and would not have been motivated to add additional, unnecessary circuitry to Ellsberry. For this additional reason, a POSA would not have been motivated to modify Ellsberry as Samsung proposes.




Fig. 5

EX1005, FIG. 5 (annotated).

IX. CONCLUSION

In signing this declaration, I recognize that the declaration will be filed as evidence in a post-grant review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required, I will appear for cross-

examination within the United States during the time allotted.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed on this 2nd day of June 2026, in Ann Arbor, Michigan.

A handwritten signature in black ink, appearing to read 'Steven Przybylski', written in a cursive style.

Steven Przybylski