

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

QUALCOMM INCORPORATED and
SAMSUNG ELCETRONICS CO., LTD.,

Petitioner,

v.

NETWORK SYSTEM TECHNOLOGIES, LLC,

Patent Owner.

IPR2024-00343
U.S. Patent No. 7,594,052

**EXPERT DECLARATION OF MARWAN HASSOUN IN SUPPORT OF
PETITIONER REPLY**

TABLE OF CONTENTS

I.	Introduction.....	1
II.	Qualifications.....	1
III.	Materials Considered.....	2
IV.	Relevant Legal Principles.....	2
V.	Claim Construction.....	2
	A. Bus, Switch, Point-to-Point Wire, and Network are All Identified as Example Structure for “Interconnect Means”	2
	B. Multi-Hop Interconnection Between Modules Separated by One or More Network Nodes.....	8
VI.	Responsive Opinions	12
	A. Ground 1: Anticipation by Radulescu2004.....	12
	1. Radulescu2004 Discloses That the Communication Service Request is Based On an Address Range	12
	2. Radulescu2004 Discloses “Mapping” as Claimed	21
	B. Ground 2: Radulescu2004 in view of OCP.....	31
	1. A POSA Would Have Been Motivated in View of OCP to Use MAddrSpace as a Communication Service Identification	31

2.	A POSA Would Have Been Motivated in View of OCP to Use Threads as a Communication Service Identification	34
3.	OCP and Radulescu2004’s Connections are Compatible	38
4.	Any Deficiencies in Radulescu2004 Regarding Mapping Are Cured by OCP	40
C.	Ground 3: Anticipation by Weber	41
1.	All Embodiments of Weber disclose a Single Integrated Circuit	41
2.	Weber has Same Structure as ‘052 for Mapping Means.....	48
3.	Weber Discloses an Address Range	53
4.	Weber Discloses an Interconnect Means Under All Constructions	56
5.	Weber’s Bus-with-Interfaces is a “Network”	68
6.	Weber’s Interfaces are “Network” Interfaces	72
D.	Ground 4: Weber and Radulescu2004.....	72
1.	It Would Have Been Obvious to Implement Weber- Radulescu2004 Combination on a Single Integrated Circuit.....	72
2.	The Reasons Provided in My First Declaration Do Provide Motivation to Combine Embodiments of Weber	73

3.	Sufficient Detail was Provided for a POSA to Combine Weber and Radulescu2004.....	74
E.	IEEE	76

I, Marwan Hassoun, declare as follows:

I. INTRODUCTION

1. I have been retained by Petitioner Qualcomm Incorporated as an independent expert consultant in this proceeding before the United States Patent and Trademark Office. I understand that this declaration involves my expert opinion and expert knowledge related to U.S. Patent No. 7,594,052 (the “‘052 patent”), titled “Integrated Circuit and Method of Communication Service Mapping,” and its field of endeavor.

2. I have reviewed the response of Patent Owner (Paper 38) and the Expert Declaration of Dr. Nader Bagherzadeh (Ex. 2070). The statements made herein are in response and are based on my own knowledge and opinions. I did not find the arguments in the Patent Owner response and Dr. Bagherzadeh’s Declaration persuasive, and the arguments did not change my opinion that the challenged claims are invalid as I explained in my first declaration (Ex. 1003).

II. QUALIFICATIONS

3. My qualifications and background are set forth in my first Declaration in Support of Petition for *Inter Partes* Review of U.S. Patent No. 7,594,052 (Ex. 1003) (“First Declaration”). My curriculum vitae has been included as Ex. 1004.

4. As in my First Declaration, my opinions in this declaration are based on the perspective of a person of ordinary skill in the art in the early 2004 timeframe. This is true even if my statements are in the present tense.

III. MATERIALS CONSIDERED

5. My technical review, analysis, insights, and opinions are based on my experience and other qualifications discussed above, as well as my study of relevant materials.

6. I have reviewed and am familiar with the materials discussed in my First Declaration, as well as any additional materials discussed in this reply declaration.

7. I have also reviewed and am familiar with Dr. Bagherzadeh's Declaration (Ex. 2070), as well as the exhibits cited therein.

IV. RELEVANT LEGAL PRINCIPLES

8. My understanding of the relevant legal principles is set forth in my First Declaration.

V. CLAIM CONSTRUCTION

A. Bus, Switch, Point-to-Point Wire, and Network are All Identified as Example Structure for "Interconnect Means"

9. I do not believe the term "interconnect means" in the claims of the '052 patent is limited to a network. The general understanding in the art at the time of the '052 was that the term "interconnect" was understood to broadly refer to any

structure used for interconnecting modules. For example, the IEEE Dictionary of 2000 defines interconnect as 1) “[a] collective term for structures (in an integrated circuit) that propagate a signal between the pins of cell instances with as little change as possible.” Ex. 1052, 4. This definition does not exclude a bus from being an “interconnect” or limit an interconnect to being a network. Dr. Bagherazadeh also agreed that the common understanding of “interconnect” would include things like a bus. Ex. 1050, 11:21-22 (“Buses are a form of an interconnection. I’m not disputing that.”); Ex. 2070, ¶112 (“conventional interconnects like buses or switches”).

10. I do not believe a POSA would understand the ‘052 specification or claims to be redefining the meaning of the term “interconnect.” Rather, I believe a POSA would understand the ‘052 specification and claims to be using “interconnect” in the manner consistent with the general understanding of the term. As discussed in my first declaration, the ‘052 patent clearly identifies each of a bus, switch, point-to-point wire, and network as example structures of an interconnect, and this aligns with the general understanding that the term broadly includes any structure that interconnects modules together. Ex. 1003, ¶82.

11. Dr. Bagherzadeh bases his opinion that “interconnect means” is limited to a network on his assertion that the specification “distinguishes the claimed interconnect or network from conventional interconnects like buses or

switches.” Ex. 2070, ¶¶112. The specification, however, does not distinguish an “interconnect” from a direct bus, the specification distinguishes a “network” or “NoC” from a direct bus. Ex. 1001, 1:48-54 (“This is because, in contrast to existing on-chip interconnects (e.g., buses, switches, or point-to-point wires), where the communicating modules are directly connected, in a NoC the modules communicate remotely via network nodes.”)¹; 2:31-34 (“Introducing networks as on-chip interconnects radically changes the communication when compared to direct interconnects such as buses or switches.”); 2:34-42 (“This is because of the multi-hop nature of a network, where communication modules are not directly connected, but separated by one or more network nodes.”). Since these statements distinguish a “network” or “NoC,” not an “interconnect” more broadly, a POSA would not understand these statements to limit the term “interconnect means” in the claims.

12. Dr. Bagherzadeh also bases his opinion that “interconnect means” is limited to a network on his assertion the specification uses the terms “interconnect means” and “network” interchangeably. Ex. 2070, ¶¶113. The patent, however, does not use the terms interchangeably. The patent consistently uses the word “network” when it is referring to a network and the word “interconnect” when it is

¹ All emphasis is added unless otherwise noted.

referring to something more broadly – something that is inclusive of buses, point-to-point wires, and networks. *See e.g.*, Ex. 1001, 1:33-2:53.

13. A POSA would understand the patent to be describing a network as *an example* of an interconnect or interconnect means, not the only possible interconnect or interconnect means. Dr. Bagherzadeh references the “N” in the claims to support his assertion that “interconnect means” and “network” are used interchangeably. Ex. 2070, ¶¶104, 107. I do not believe a POSA would understand the “N” in the claims to overrule the understanding of “interconnect” known to those in the art and as used in the description of the ‘052 patent. Instead, the “N” would be understood to be a reference to the *example* interconnects shown in the Figures of the patent, which coincides with the manner in which “N” is used throughout the ‘052 patent. *See* Ex. 1001, 5:6-23, FIGs. 1-4.

14. Dr. Bagherzadeh also points to the sentence in the description that reads “the interconnect is embodied as a network on chip NOC, allowing communication between all aforementioned modules.” Ex. 2070, ¶109 (citing Ex. 1001, 4:46-50). However, this sentence again indicates that a network is *an example* of (“is embodied as”) an interconnect. This sentence does not state that the term interconnect *excludes* other types of structures. Since the patent does not equate “interconnect” with “network” as suggested by Dr. Bagherzadeh, this also is

not a reason that a POSA would understand the term “interconnect means” to be limited to a network.

15. Based on all of this I believe a POSA would understand that the term “interconnect means” to correspond all the structure for an “interconnect” and not be limited to a network. A POSA would understand that if the inventors had truly wanted to limit the “interconnect means” clause to a network, they would have used the word “network” or “network means,” not “interconnect means.”

16. Moreover, Dr. Bagherzadeh did not know that the specification described bus, switch, and point-to-point wires as examples of an interconnect. Ex. 1050, 8:14-10:24 (Reproduced below.)

14 Q. I'm going to be asking some of the meanings of
15 those words, so this is going to be important. Let me
16 start with this -- this term "interconnect."
17 Now, interconnect appears in the claims of the
18 '052 patent in connection with something called an
19 interconnect means; right?
20 A. Interconnect appears in '052 patent. That's
21 correct.
22 Q. And in the claims as well?
23 A. Yes. I believe so, yeah.
24 Q. Okay. The patent uses the term "interconnect"
25 to describe point-to-point wires, buses, and switches;

Page 9

1 correct?
2 MR. CHRISTOPHER: Objection to form.
3 THE WITNESS: No, it doesn't.
4 BY MR. VANDENBURGH:
5 Q. I'm talking about the word "interconnect"
6 alone.
7 In your opinion, is -- the patent doesn't use
8 interconnect to describe a point-to-point wire, bus, or
9 switch?
10 MR. CHRISTOPHER: Objection to form.
11 THE WITNESS: The patent is very clear. It's
12 referring to multi-hop networks with one or more nodes.
13 BY MR. VANDENBURGH:
14 Q. I understand that that might be your opinion
15 with respect to interconnect means, but I'm only asking
16 you about the word "interconnect" right now.
17 Your testimony is that the patent doesn't use
18 interconnect to describe wires, switches, or buses?
19 MR. CHRISTOPHER: Objection to form.
20 THE WITNESS: A POSITA reading '052 will
21 understand interconnect that they're referring to is not
22 bus, wires, or point-to-point.
23 BY MR. VANDENBURGH:
24 Q. Okay. Let me get more specific. Can I direct
25 your attention to Column 1, sentence beginning at

1 Line 50.
 2 A. I'm sorry. Counsel, you're referring to the
 3 patent; right?
 4 MR. VANDENBURGH: Correct. So I'm in -- and
 5 that's a good point.
 6 For the record, I'm referring to Exhibit 1001,
 7 which is the '052 patent, Column 1, Line 50.
 8 A. Okay. Yes.
 9 Q. I'm going to read it for the record.
 10 "This is because, in contrast to existing
 11 on-chip interconnects (e.g., buses, switches, or
 12 point-to-point wires), where the communicating modules
 13 are directly connected, in a NoC the modules communicate
 14 remotely via network nodes."
 15 First of all, did I read that correctly?
 16 A. Yes, you did.
 17 Q. And it is true, isn't it, that that statement
 18 refers to buses, switches, or point-to-point wires as
 19 examples of existing on-chip interconnects?
 20 A. It -- it is just giving you an example, but the
 21 interconnect that is mentioned throughout this '052 is
 22 all about a NoC. You could say that you could connect
 23 any way you want, but '052 will come out for a POSITA
 24 that it's -- it is a Network-on-Chip.

B. Multi-Hop Interconnection Between Modules Separated by One or More Network Nodes

17. I understand Dr. Bagherzadeh's opinion to be that regardless of whether the term "interconnect means" is a means-plus-function term, the term is limited to a network and his proposed construction is "multi-hop interconnection between modules separated by one or more network nodes." Ex. 2070, ¶¶100. A POSA would understand "multi-hop" and "separated by one or more network nodes" as being synonymous. In other words, a "multi-hop interconnection between modules" *is the same thing as* "an interconnection between modules separated by one or more network nodes." Ex. 1001, 2:32-37("Introducing

networks as on-chip interconnects radically *changes the communication* when compared to *direct interconnects*, such as buses or switches. This is because of the *multi-hop nature of a network*, where communication modules are not directly connected, but *separated* by *one* or more *network nodes*.”). Because the two phrases mean the same thing, I do not believe a POSA would understand a network in ‘052 to be *both* multi-hop *and* separated by one or more network nodes. Rather, a POSA would understand the ‘052 to describe a network as “a multi-hop interconnection between modules” or alternatively as “an interconnection between modules separated by one or more network nodes.”

18. From the general understanding in the art, a POSA would understand that a “hop” is a communication link with a node on each end. Webster’s Telecom Dictionary of 2008 defines hop in a similar manner. Ex. 1051, 6 (“In networks, the journey a signal makes across a transmission link between two devices such as bridges, hubs, switches, or routers.”). Each link in an interconnect having a node on each end is a hop. When data traverses from one node across a link to another node, it is said to have “hopped.” Thus, each node of an interconnect creates an additional hop from a master to a slave. If there are zero nodes between a master and a slave, i.e. the master and slave are directly connected, there is a single hop. If there is one node between a master and a slave, there are two hops – one hop from

the master to the intermediate node and a second hop from the intermediate node to the slave. Two nodes between master and slave creates three hops and so forth.

19. The patent uses the term “hop” in the same manner. The patent uses the term “hop” just once, in the term “multi-hop.” Ex. 1001, 2:34-36. The patent describes what “multi-hop” means; it means the modules are “separated by one or more network nodes.” *Id.* (“because of the multi-hop nature of a network, where communication modules are not directly connected, but separated by one or more network nodes”). Thus, “multi-hop” and “separated by one or more network nodes” refer to the same thing. Both phrases refer to the fact that the interconnect has one or more network nodes between the modules. As discussed above, each node between the modules creates an additional hop making the interconnect “multi-hop.” One or more nodes between the modules also separates the modules making the modules “separated by one or more network nodes.” This sentence also suggests that whenever modules are not “directly connected,” the path is “multi-hop,” which is also consistent with the understanding that each node between the modules creates a hop.

20. Dr. Bagherzadeh argues: “a ‘hop’ is a router/switch that a packet originating from an arbitrary protocol may traverse.” Ex. 2070, ¶114. I do not believe a POSA would understand “hop” to be limited to a router or switch. Webster’s Telecom Dictionary of 2008 defines hop in a non-limiting manner as

between “two devices” and specifically lists devices other than switches and routers as examples of such devices. Ex. 1051, 6 (“In networks, the journey a signal makes across a transmission link between two devices such as bridges, hubs, switches, or routers.”) (emphasis added). The Data & Computer Communications textbook submitted by Patent Owner also indicates that a hop is not limited to switches and routers by indicating that a bridge provides a hop. Ex. 2010, 461 (“It would appear preferable for bridge 107 to forward the frame, since it will involve only one hop, whereas if the frame travels through the bridge 101, it must suffer two hops.”). The patent also does not limit hop to only switches and routers. The reference in the patent to a network being “multi-hop,” suggests that it includes one or more nodes and associated links, but do not put specific restrictions on those nodes or the links. Ex. 1001, 2:33-36.

21. The patent also does not limit a “network node” to being only a switch or router. The patent also does not provide a specific definition of “network node” vs. simply “node,” suggesting that “network node” merely refers to any node in a network, which is any node between the modules. The patent uses the term “network node” exactly twice and both uses suggest that any node between the modules is a “network node.” Ex. 1001, 1:52-53 (“where the communicating modules are directly connected, in a NoC the modules communicate remotely via network nodes.”), 2:33-37 (“where communication modules are not directly

connected, but separated by one or more network nodes”). Thus, a POSA would understand that neither “hop” nor “network node” require that switch and router be used.

VI. RESPONSIVE OPINIONS

A. Ground 1: Anticipation by Radulescu2004

1. Radulescu2004 Discloses That the Communication Service Request is Based On an Address Range

22. Dr. Bagherzadeh argues that claim element [6.pre.2] (and similar limitations in claims 1 and 7) require the communication service request to “comprise” (i.e., contain) and an address range or thread. Ex. 2070, ¶¶146-161; e.g., ¶154 (“a POSITA would understand that an address range is *not communicated in a request* if it ‘is configurable in the narrowcast module.’”); ¶159 (“In other words, because Rădulescu2004 does not teach that the address range as a communication service identification *is included in the request issued* by the master module (M), Rădulescu2004 does not disclose the ‘address range’ limitation.”); ¶150 (“a POSITA would understand that Rădulescu2004 issues a request comprising an ‘address’ to identify each slave and must issue multiple requests to identify multiple locations within the same slave. Nothing in Rădulescu2004 suggests a communication service identification that is an address range.”). I disagree with this because that is not what the claim says or how a POSA would understand the claim.

23. Claim 6 only requires that requests be “based on” a communication service identification that is either an address range or thread. Ex. 1001, 10:9-15:

6. Method of communication service mapping in an integrated circuit, having a plurality of processing modules (M, S), wherein at least one first of said processing modules (M) requests at least one communication service to at least one second processing module (S) **based on specific communication properties and at least one communication service identification**, wherein said at least one communication service identification comprises at least one **communication thread or at least one address range**, said address range for identifying one or more second processing modules (S) or a memory region within said one or more second processing modules (S), comprising the steps of:

Dr. Bagherzadeh appears to interpret “based on” as *contains* (i.e. “included in the request issues” or “communicated in a request” (Ex. 2070, ¶154, ¶159)). This imports an additional requirement that is not in the plain language of the claim.

24. Dr. Bagherzadeh’s interpretation of “based on” as *contains* is inconsistent with how the specification of the ‘052 patent describes the purported invention. I have reviewed the specification of the ‘052 patent, and it never discloses that the service request itself contains an address range (i.e. “included in the request issues” or “communicated in a request” (Ex. 2070, ¶154, ¶159)). The ‘052 patent states that the “invention is based ... mapping the identification means ... to protocols” and that the “identification means” are “communication threads and addresses.” Ex. 1001, 4:15-21:

The invention is based on the idea to offer differentiated services for protocols such as DTL, MTL, AXI, and OCP, by mapping the identification means of these protocols to connections. The identification means in the existing protocols are: communication threads and addresses, i.e. the threads or addresses are mapped to connections through the interconnect based on specific connection properties.

25. The '052 patent specification goes on to consistently describe that addresses—not address ranges—are included in the requests and those transmitted addresses are used in the mapping process. Ex. 1001, 6:48-53 (“The basic operation of such a shell associated to a master as shown in FIG. 2A is to sequentialize commands cmd and their flags flag, *addresses adr*, and write data wr_data in request messages msg, and to desequenzialize messages msg into read data rd_data, and write responses wr_resp.”); 5:61-6:3 (“In other words, the adapter units A translate the requested communication service into messages msg, which are processed by the network interface NI and are translated into flits fit to be transmitted over the network N, in particular the adaptor unit A maps the identification means of these protocols to connections. The identification means in the existing protocols are: communication threads and *addresses*, i.e. the threads or *addresses are mapped to connections* through the interconnect based on specific connection properties.”); 4:18-21 (“The identification means in the existing protocols are: communication threads and addresses, i.e. the threads or *addresses are mapped to connections* through the interconnect based on specific connection properties.”). This is consistent with the depiction in Fig. 2A. Fig. 2A shows the

“addr” signal (not an address range), being input from the master into the adapter unit:




FIG.2A

Ex. 1001, 6:48-53, Fig. 2A (annotated). Dr. Bagherzadeh, in reference to the nearly identical figure in Radulescu2004 (reproduced below), testified that a “POSITA would immediately recognize that ‘addr’ represents an address, and not an address range.” Ex. 2070, ¶156 (“I note that Rădulescu2004 uses the term ‘addr’ when referencing the address. Ex. 1005 at FIGs. 5, 6. *A POSITA would immediately recognize that ‘addr’ represents an address, and not an address range.* The same authors of Rădulescu2004 similarly used to the term ‘address (ADDR)’ to represent an address in a later version of the same article. Ex. 2078 at 8. This is consistent with how a POSITA would understand Rădulescu2004’s recitation of ‘addr’ as an address, and not an address range. Similarly, the same authors of

Rădulescu2004 used to term ‘ADDR_SPC’ to represent an address space, which a POSITA would understand is analogous to an address range. Ex. 1002 at 91, 94. This is consistent with how a POSITA would understand the ‘052 Patent’s recitation of an address range. Rădulescu2004’s prior use of the term ‘ADDR’ as being different from an ‘ADDR_SPC’ further supports a POSITA’s undersatnding [sic] Rădulescu2004 discloses an address (‘addr’) and not an address range.”).




Figure 5. Master shell

Ex. 1005, Fig. 5. Dr. Bagherzadeh has not pointed to any disclosure in the ‘052 patent that indicates that the service request from the master itself contains an address range.

26. The ‘052 patent specification discusses “address ranges” in connection with Figure 4. Ex. 1001, 8:7-24:

FIG. 4 shows a basic block diagram of a network on chip for illustrating the operation of the mapping of IP block communication via address ranges to connections. The network on chip comprises a master M and two slaves S1, S2 each being connected to a network via a network interface NI, respectively. The two slaves S1, S2 are implemented as memories, wherein the first slave S1 comprises the address range 0000-7FFF and the second slave S2 comprises the address range 8000-FFFF. Accordingly, the network interface NI associated to the master M may map the communication to the first slave S1, i.e. to the address range 0000-7FFF, to a connection with the connection id cid=0 and the communication to the second slave S2, i.e. to the address range 8000-FFFF, to a connection with the connection id cid=1.

Nothing in this portion of the specification says that the service request sent by a module (masters or slaves) comprises or contains an address range. Instead, this section of the specification discusses, and Figure 4 shows, a table of address ranges with their associated connections are stored within the network interface NI:




FIG.4

Ex. 1001, Fig. 4 (annotated). The specification teaches that the service requests from master M containing “identification means” such as “threads or addresses”

are sent to an adapter unit within the network interface NI. Ex. 1001, 5:61-6:3 (“In other words, the adapter units A translate the requested communication service into messages msg, which are processed by the network interface NI and are translated into flits fit to be transmitted over the network N, *in particular the adaptor unit A maps the identification means of these protocols to connections*. The identification means in the existing protocols are: *communication threads and addresses, i.e. the threads or addresses are mapped to connections* through the interconnect based on specific connection properties.”); 4:18-21 (“The identification means in the existing protocols are: *communication threads and addresses, i.e. the threads or addresses are mapped to connections* through the interconnect based on specific connection properties.”). The ‘052 patent specification teaches and a POSA would understand that for address range-based mapping, the network interface uses the address received in the request and the address ranges within the network interface to map the request to a connection (e.g. “The address generation and synchronization is performed in the shell. Each stream of data has a predefined buffer in the memory, and, hence, *a predefined address range associated in the shell*. Consequently, *based on the addresses* generated by the shell, a network interface can identify each stream, and map it to the right connection having connection properties,” and as depicted in FIG. 4 above). Ex. 1001, 8:7-36:

FIG. 4 shows a basic block diagram of a network on chip for illustrating the operation of the mapping of IP block communication via address ranges to connections. The network on chip comprises a master M and two slaves S1, S2 each being connected to a network via a network interface NI, respectively. The two slaves S1, S2 are implemented as memories, wherein the first slave S1 comprises the address range 0000-7FFF and the second slave S2 comprises the address range 8000-FFFF. Accordingly, the network interface NI associated to the master M may map the communication to the first slave S1, i.e. to the address range 0000-7FFF, to a connection with the connection id cid=0 and the communication to the second slave S2, i.e. to the address range 8000-FFFF, to a connection with the connection id cid=1.

Hence, protocols where there are no communication threads (e.g., DTL, MTL), address ranges can be mapped to specific connections. An address range can identify one or more slave modules, or a memory region within a slave module, e.g. a buffer used in some particular communication.

The above described mapping of address ranges to a connection can for example be applied to an Eclipse shell. Here, streams of data between modules are mapped on memory buffers. Producer and consumer modules read and write the streams via a shell using an identifier. The address generation and synchronization is performed in the shell. Each stream of data has a predefined buffer in the memory, and, hence, a predefined address range associated in the shell. Consequently, based on the addresses generated by the shell, a network interface can identify each stream, and map it to the right connection having connection properties.

For example, a POSA would understand that a request with address 0004 (which is within the range 0000-7FFF assigned to slave S1) would be mapped by the network interface to connection id cid=0, and a request with address 8FFF (which is within the range 8000-BFFF²) would be mapped by the network interface to

² As I noted in my First Declaration, the specification and Fig. 4 depict different implementations for S2. The specification states that “The two slaves S1, S2 are implemented as memories, wherein the first slave S1 comprises the address range 0000-7FFF and the second slave S2 comprises the address range 8000-FFFF.

connection id cid=1. A POSA would understand that the service request sent by master M is “based on” an address range in the sense that the request contains an address that corresponds to one of the address ranges stored in the network interface that identifies the slave S for which the message is intended. Ex. 1001, 8:7-20; *see also* 8:21-40 (reproduced above).

27. A POSA would understand that Radulescu2004 works in the same way as I just described. As I explained in my First Declaration, the address range assigned to each slave is stored in the narrowcast shell within the network interface. Ex. 1003, ¶124; Ex. 1005, 4 (“Within a narrowcast connection, the slave for which the transaction is destined is selected based on the address (Conn block). *The address range assigned to a slave is configurable* in the narrowcast module.”). Just like in the ‘052 patent, the communication service requests in

Accordingly, the network interface NI associated to the master M may map the communication to the first slave S1, i.e. to the address range 0000-7FFF, to a connection with the connection id cid=0 and the *communication to the second slave S2, i.e. to the address range 8000-FFFF, to a connection with the connection id cid=1.*” Ex. 1001, 8:15-20. Fig. 4 appears to indicate that connection cid=1 is used only for the address range 8000-BFFF, while a third connection cid=2 is used for the address range C000-FFFF. Ex. 1001, Fig.4.

Radulescu2004 are “based on” these address ranges that (again, just like in the ‘052 patent FIG. 4 above and described at least at 8:30-36) are in the Network Interface (NI) because the requests contain an address that, using the address ranges stored in the narrowcast shell, can be used to identify “the slave for which the transaction is destined.” Ex. 1005, 4; Fig. 5; Ex. 1001, Fig. 4, 8:31-36 (“Each stream of data has a predefined buffer in the memory, and, hence, *a predefined address range associated in the shell*. Consequently, based on the addresses generated by the shell, a network interface can identify each stream, and *map it to the right connection having connection properties*.”). Neither of these passages state that the network is considering the connection properties at the time a connection is being chosen; rather, just as in Radulescu2004, they are stating that addresses or threads are used to choose the appropriate connection and its associated connection properties. A POSA would understand that this is all that is required by limitation [6.pre.2] (and similar limitations in claims 1 and 7).

2. Radulescu2004 Discloses “Mapping” as Claimed

28. Dr. Bagherzadeh argues that “Radulescu2004 does not disclose ‘mapping the requested at least one communication service based on said specific communication properties to a connection based on a set of connection properties according to at least one communication service identification,’ as recited in claim 6 of the ‘052 Patent because Radulescu2004 does not disclose the claimed

‘communication service identification.’” Ex. 2070, ¶162. Dr. Bagherzadeh is wrong. As I discussed in the proceeding section, Radulescu2004 does disclose a communication service identification as claimed and discloses mapping according to that communication service identification.

29. Dr. Bagherzadeh further argues that “Radulescu2004 does not satisfy the ‘mapping’ step claimed in the ‘052 Patent,” and “[i]n particular, Radulescu2004 lacks the claimed ‘mapping ... to a connection based on a set of connection properties.” Ex. 2070, ¶¶164-165. Dr. Bagherzadeh is incorrect.

30. Dr. Bagherzadeh does not dispute that Radulescu2004 discloses the use of connections, each have a set of connection properties. Ex. 2070, ¶92:

Radulescu2004 describes ... ‘(b) the NI shells, *which implement the connections* (e.g., narrowcast, multicast), transaction ordering for connections, and other higher-level issues specific to the protocol offered to the IP.’ ... The NI kernel ‘communicates with the NI shells via ports,’ where each ‘port can have *multiple connections to allow differentiated traffic classes* in which case there are also CONNID signals to select on which connection a message is supplied or consumed.’

See also Ex. 2070, §IX.A (not disputing element [6.b]). These connections allow for “differentiated traffic classes” by “select[ing] on which connection a message is supplied or consumed.” Ex. 1005, 3 (“A port can have multiple connections to allow differentiated traffic classes, in which case there are also connid signals to select on which connection a message is supplied or consumed.”); 2 (“Connections are composed of unidirectional point-to-point channels (between a single master

and a single slave). To each channel, properties are attached, such as guaranteed message delivery or not, in order or un-ordered message delivery, and with or without timing guarantees.”); Abstract (“It provides both guaranteed and best-effort services via connections.”); 1 (“We provide transport-layer services by defining connections (e.g., point-to-point or multicast) configured for specific properties (e.g., throughput, ordering).”); 6 (“We offer, via connections, high-level services, such as transaction ordering, throughput and latency guarantees, and end-to-end flow control.”). A POSA would understand that by offering differentiated services (e.g., Ex. 1005, 1 (“provides both guaranteed and best effort services via connections” ... “provide transport-layer services by defining connections ... configured for specific properties (e.g., throughput, ordering)”)), 3 (“multiple connections to allow differentiated traffic classes, in which case there are also connid signals to select on which connection a message is supplied or consumed.”))—that Radulescu2004 is mapping to specific connections based on connection properties, and is doing it “according to” the service identification of an address range. As I explained in my First Declaration, each connection has connection properties. Ex. 1003, ¶¶131-133, 139. Thus, the mapping of a service request to one of the point-to-point connections in the narrowcast shell is to “a connection based on a set of connection properties” as the claim required. Ex. 1005, 3, 4, 1, 2; Ex. 1003, ¶¶140-143.

31. Dr. Bagherzadeh argues that “Radulescu does not consider the ‘connection data’ when selecting the connection for the narrowcast transmission” and “merely discloses selecting the slave ‘based on the address provided in the transaction’ without regard to the ‘particular set of transaction requirements.’” Ex. 2071, ¶170. Dr. Bagherzadeh’s argument is a repackaging of an argument that was preliminarily rejected by the Board in its Institution Decision. Patent Owner argued that the claims required that connection properties be assigned at the time service requests are mapped by the network interface and excluded a fixed hardware configuration where properties are assigned “prior to their use.” Paper 7, 46-50, *see e.g.*, 47 (“Rădulescu2004 does not disclose a communication service mapping as recited in claim 6 of the ’052 Patent. ... Instead, Rădulescu2004 discloses **a fixed hardware configuration** between modules and network interfaces that is required to be set up before receiving service requests during an operational stage, i.e., **prior to mapping communication service requests.**”); 48 (“The specification and figures of the ’052 Patent further highlight the distinction between the **creation of desired properties prior to their use**, as taught in Rădulescu2004, and the claimed invention.”); 50 (“Rădulescu2004 discloses the **creation of desired properties prior to their use**, which is different from the communication service mapping recited in claim 6.”). The Board properly rejected this argument, finding no basis in

the claim language to exclude such a fixed configuration. Paper 11, 21-24, 29-30;

e.g.:

Patent Owner argues that Rădulescu2004 does not disclose the recited communication service mapping. Prelim. Resp. 47–50. Patent Owner asserts that, instead, Rădulescu2004 discloses a fixed hardware configuration between modules and network interfaces that is required to be set up before receiving service requests during an operational stage, i.e., ***prior to mapping communication service requests***. *Id.* at 47. Patent Owner argues that the recited mapping communication service requests excludes such a hardware configuration. *Id.* at 48.

We disagree with this argument by Patent Owner. The plain language of the term provides no such exclusion. Ex. 1001, 10:7–32. Patent Owner cites nothing from the Specification that defines communication service requests to have such an exclusion. Prelim. Resp. 47–50. Although Patent Owner cites lines 7–26 of column 8 of the '052 patent, ***that disclosure does not provide a definition for communication service requests that excludes a fixed hardware configuration.*** Prelim Resp. 49; Ex. 1001, 8:7–26.

Further, ***during prosecution of the '052 patent***, the Applicant ***attempted to distinguish*** an earlier version of issued claim 6 that recited 'communication service mapping' from the cited references ***Dielissen and Rădulescu2000***. Ex. 1002, 106, 110. The applicant argued that claim 6 was distinguishable ***because Dielissen and Rădulescu2000 had a fixed hardware configuration***. *Id.* at 110. The Examiner found the Applicant's argument to be unpersuasive, and maintained the rejection of claim 6 and the other pending claims. *Id.* at 119, 122. The Examiner indicated that the distinction the Applicant was raising ***was not recited in the claims*** ... It was not until the Applicant ***added the communication thread/address range limitation*** to the pending claims that the Examiner allowed claim 6 and the other pending application claims to issue.

Paper 11, 22-23.

Patent Owner argues that Rădulescu2004 does not disclose communication service mapping as recited in claim 6 of the '052 patent, but discloses a fixed hardware configuration between modules

and network interfaces that is required to be set up prior to receiving service requests during an operational stage, i.e., ***prior to mapping communication service requests***. Prelim. Resp. 52. As set forth in Section III.A.2.b above, ***we disagree with Patent Owner’s argument that claim 6 distinguishes a fixed hardware configuration***. After reviewing the record, we determine that Petitioner has made a sufficient showing for limitation 6.d.

Paper 11, 30. Dr. Bagherzadeh’s argument seeks to impose the same temporal restriction that the Board properly rejected.

32. The claim recites “mapping ... to a connection based on a set of connection properties according to said at least one communication service identification.” Ex. 1001, cl. 6. The mapping of the service request is “to a connection based on a set of connection properties” even if those connection properties were assigned to the connection prior to mapping. *Id.* A POSA would understand that the plain reading of “a connection based on a set of connection properties” in this element means a connection has a set of connection properties. The phrase “based on a set of connection properties” directly follows after “a connection,” which indicates that the connection is what is “based on a set of connection properties.” This is also consistent with the specification, which teaches that “[a] connection is considered as a set of channels, ***each having a set of connection properties....***” Ex. 1001, 4:52-54. A POSA would understand that this

part of element [6.d] is met by mapping to a connection that has a set of connection properties.³

33. My understanding of this claim element is supported by the prosecution history of the ‘052 patent, which found this limitation met by similar inventor publications describing the Aethereal NoC. During prosecution of the ‘052 patent, the Examiner found that the Dielissen (Ex. 1011) and Radulescu2000 (Ex. 1010) publication disclosed the “to a connection based on a set of connection properties” limitation. Ex. 1002, 79-82, 119-123; Ex. 1011, §2.2.1, §2.2.2, Figs. 1, 7; Ex. 1010 at 14. The Applicant did not dispute this. The disclosures in Dielissen and Radulescu2000 are very similar to the disclosure in Radulescu2004 with respect to this issue. Each of Dielissen, Radulescu2000, and Radulescu2004 describe the Aethereal NoC. Dielissen Sections 2.2.1 and 2.2.2 are particularly similar to Radulescu2004 Sections 4.1 and 4.2. The primary difference that I see is that Radulescu2004 adds disclosure. For example, Radulescu2004 adds disclosure regarding the narrowcast shell/narrowcast connections and discusses “selection

³ I will also note that Dr. Bagherzadeh incorrectly states that “the only ‘connection property’ disclosed in Radulescu2004 is ordering....” Ex. 2070, ¶167. As I discussed in my First Declaration, Radulescu2004 teaches additional connection properties such as throughput and latency. Ex. 1003, ¶131.

based on addresses” and that the “address range assigned to a slave is configurable in the narrowcast module.” *Compare* Ex. 1005, §§4.1, 4.2 with Ex. 1011, §§2.2.1, 2.2.2.

34. As I discussed in my First Declaration, Radulescu2004 shows the mapping as claimed (Ex. 1003, ¶¶139-143, 130-133) and in the same way as disclosed by the ‘052 patent specification. For example, the ‘052 patent describes using an address to determine which connection to send a request over where connection properties are not directly used. Ex. 1003, ¶¶141-142; Ex. 1001 at 8:7-20:

FIG. 4 shows a basic block diagram of a network on chip for illustrating the operation of the mapping of IP block communication via address ranges to connections. The network on chip comprises a master M and two slaves S1, S2 each being connected to a network via a network interface NI, respectively. The two slaves S1, S2 are implemented as memories, wherein the first slave S1 comprises the address range 0000-7FFF and the second slave S2 comprises the address range 8000-FFFF. Accordingly, the network interface NI associated to the master M may map the communication to the first slave S1, i.e. to the address range 0000-7FFF, to a connection with the connection id cid=0 and the communication to the second slave S2, i.e. to the address range 8000-FFFF, to a connection with the connection id cid=1.

For example, Figure 4 (which Dr. Bagherzadeh argues shows the mapping step (Ex. 2070, ¶88)) shows a map of address ranges and corresponding connection stored in the network interface:




FIG.4

Ex. 1001, Fig. 4 (annotated).

35. The '052 patent describes that mapping via address ranges is done based on an association between address ranges and connection, and it does not mention any connection properties. Ex. 1001, 8:7-20 (reproduced above); Ex. 1003, ¶¶64-66, 141-142. Radulescu2004 teaches the same thing. Ex. 1005, 4 (“Narrow connections are *connections between one master and several slaves*, where each transaction is executed by a single slave selected based on the address provided in the transaction ... We implement the narrowcast connection as *a collection of point-to-point connections, one for each master-slave pair*. Within a *narrowcast connection, the slave for which the transaction is destined is selected based on the address (Conn block)*. The *address range assigned to a slave is configurable in the narrowcast module*. To provide *in-order response delivery*,

the narrowcast must also keep a history of connection identifiers of the transactions including responses (e.g., reads, and acknowledged writes), and the length of these responses.”); Ex. 1003, ¶¶140-143.

36. A POSA would understand that connection properties are taken into account during creation of the connections in both the ‘052 patent and Radulescu2004, and that the mapping to the connection (that has already been created) is therefore “based on a set of connection properties.” Ex. 1001, Ex. 1001, 5:43-52:

Connections must be created stating the requested service level. The request for the connection are either accepted or rejected. Connection acceptance may lead to resource reservations in the network, e. g. buffers or a link bandwidth percentage. After usage, the connection is closed and these resources are freed. Different connections are created and closed independently, possibly at different points in time. The configurations thereof can be computed at compile time (i. e. off-line), or at run time.

Ex. 1001, 6:4-8:

Once a connection has been created, the master initiates transactions by means of requests which zero or more slaves execute, perhaps leading to a response. Examples of transactions are read, write, acknowledged write, test and set, and flush.

Ex. 1005, §1, §2; Ex. 1003, ¶¶61-63, 92, 96, 99, 131-133.

37. Dr. Bagherzadeh argues that Radulescu2004 does not disclose mapping to the “right connection” or to “a specific connection.” Ex. 2070, ¶166 (“a POSITA would understand that the claimed ‘mapping’ step requires mapping the

requested communication service to a specific connection...”). However, as described in my First Declaration, Radulescu2004 maps to “a particular one of these point-to-point connections based on the address.” Ex. 1003, ¶140. Radulescu2004 does not merely select “any available connection” as Dr. Bagherzadeh suggests. Ex. 2070, ¶169.

B. Ground 2: Radulescu2004 in view of OCP

1. A POSA Would Have Been Motivated in View of OCP to Use MAddrSpace as a Communication Service Identification

38. OCP teaches two addressing signals that can be used in combination: MAddr and MAddrSpace. MAddr identifies “the slave-dependent address of the resource targeted by the current transfer,” and MAddrSpace “is used to indicate the address region of a transfer.” Ex. 1006, 24, 26. As I explained in my First Declaration, even if the Board concludes that Radulescu2004 does not teach that the service request is “based on” a communication service identification that is an address range, a POSA would have been motivated to do so using the MAddrSpace and MAddr signals of OCP. Ex. 1003, ¶¶174-175, 177.

39. Dr. Bagherzadeh does not dispute that MAddrSpace identifies an address range. Instead, Dr. Bagherzadeh argues that I “fail to articulate why a POSITA would be motivated to use the optional MAddrSpace signal extension in Radulescu2004’s system.” Ex. 2070, ¶177. Dr. Bagherzadeh is mistaken as I did

provide such analysis in my First Declaration. My First Declaration provides that a POSA would have been motivated to achieve backwards compatibility with OCP, as expressly taught by Radulescu2004. Ex. 1003, ¶¶154-155. A POSA would know that in order to provide backward compatibility with IP modules that use the combination of MAddr and MAddrSpace, the adapter units in the network interfaces of Radulescu2004's system would need to work with both MAddr and MAddrSpace. Ex. 1006, 38 (annotated figure below indicating that MAddr and MAddrSpace are “always” part of a request if these signals are required):

Signal Groups

Some OCP fields are grouped together because they must be active at the same time. The data flow signals are divided into three signal groups: request signals, response signals, and datahandshake signals. A list of the signals that belong to each group is shown in Table 12.

Table 12 OCP Signal Groups

Group	Signal	Condition
Request	MAddr	always
	MAddrSpace	always
	MBurst	always
	MByteEn	always
	MCmd	always
	MConnID	always
	MData*	datahandshake = 0
Response	MThreadID	always
	SData	always
	SResp	always
Datahandshake	SThreadID	always
	MData*	datahandshake = 1
	MDataValid	always
	MDataThreadID	always

*MData belongs to the request group, unless the datahandshake configuration parameter is enabled. In that case it belongs to the datahandshake group.

This is because some architectures take advantage of the MAddrSpace signals to identify the “address region of a transfer” such as “the register space versus the

regular memory space.” Ex. 1006, 26. The MAddr signal would then provide the location within the “address region” in order to identify the proper data. This is consistent with the disclosure of the ’052 Patent where, for example, in reference to FIG. 4 where “An *address range can identify* one or more slave modules, or a *memory region within a slave module*, e.g. *a buffer [register]* used in some particular communication.” ’052 Patent 8:23-25.

40. For the system of Radulescu2004 to be compatible with an OCP IP module that uses MAddr and MAddrSpace, the mapping would need to be based on both addressing signals. Dr. Bagherzadeh argues that “OCP’s standard MAddr signal (i.e., an ‘address’ present in Radulescu2004) already provides ... the benefit of using MAddrSpace signal, i.e., that it would identify a system target and/or a memory region within the system target.” Ex. 2070, ¶179; *see also* ¶200. I disagree with Dr. Bagherzadeh. Dr. Bagherzadeh ignores the full disclosure of Radulescu2004. A POSA would understand that Rădulescu2004 disclosure: “Within a narrowcast connection, *the slave for which the transaction is destined is selected based on the address* (Conn block). The *address range assigned to a slave is configurable in the narrowcast module*,” (Ex. 1005, 4) to be a cohesive explanation of how the identification of the connection to the slave is identifies and the two sentences to be discussing the same process. A POSA would understand Rădulescu2004 discloses a connection to a slave is chosen via the address in the

request and such an identification of which slave such an address belongs to is through the address range of that slave that has been configured in the narrowcast module. For an OCP IP module that uses MAddr and MAddrSpace, a POSA would be motivated to use the MAddrSpace, in addition to MAddr, for mapping in order to identify the address region of the transfer. Ex. 1006, 26. A POSA would have been motivated to use both the MAddr and MAddrSpace signals for such an architecture in the mapping described by Radulescu2004.

2. A POSA Would Have Been Motivated in View of OCP to Use Threads as a Communication Service Identification

41. As I discussed in my First Declaration, a POSA would have been motivated in view of the OCP Specification (Ex. 1006) to transmit a service request based on a communication service identification that is a thread in Radulescu2004's system. Ex. 1003, ¶¶154-161, 164-173. Threading is a known functionality in on-chip communications that is known to have substantial benefits. *Id.* My First Declaration explained that a POSA would have been motivated to combine Radulescu2004 and OCP generally, as well as why and how a POSA would have been motivated to use OCP's threading functionality with Radulescu2004's system. *Id.*

42. Radulescu2004 expressly states that a primary purpose of the disclosed interconnect is to provide "backward compatibility" with "existing on-chip communication protocols" such as "OCP." Ex. 1005, 1. Dr. Bagherzadeh

argues that a POSA would not have been motivated to use OCP's threading functionality in Radulescu2004. Ex. 2070, ¶¶181-186. However, Dr. Bagherzadeh is wrong.

43. Dr. Bagherzadeh argues that the MThreadID signal “is an optional extension of the OCP protocol and is not necessarily required” “in every OCP request” and thus I failed to opine as to what would motivate a POSA “to implement threads into Radulescu2004’s system besides the mere fact that OCP teaches threads.” Ex. 2070, ¶¶181-182. First, Dr. Bagherzadeh appears to misunderstand what is meant by “optional” signals in OCP. As I said in my deposition, “a system that implements these extensions would need to have [them].” Ex. 2068, 77:1-11. A POSA would know that a module that was designed to work with MThreadID and that an interconnect designed to provide “backward compatibility” with OCP as taught by Radulescu2004, would need those signals and would need to support all signals of OCP, including MThreadID and other optional signals related to it (*see, e.g.*, annotated figure below where “Some OCP fields are grouped together because they ***must be active at the same time***,” and all the optional signals in a signal group are “always” there).

Signal Groups

Some OCP fields are grouped together because they must be active at the same time. The data flow signals are divided into three signal groups: request signals, response signals, and datahandshake signals. A list of the signals that belong to each group is shown in Table 12.

Table 12 OCP Signal Groups

Group	Signal	Condition
Request	MAddr	always
	MAddrSpace	always
	MBurst	always
	MByteEn	always
	MCmd	always
	MConnID	always
	MData*	datahandshake = 0
	MThreadID	always
Response	SData	always
	SResp	always
	SThreadID	always
Datahandshake	MData*	datahandshake = 1
	MDataValid	always
	MDataThreadID	always

*MData belongs to the request group, unless the datahandshake configuration parameter is enabled. In that case it belongs to the datahandshake group.

Ex. 1006, 38.

44. Second, in my First Declaration, I provided testimony that a POSA would be motivated to use OCP's threads. Ex. 1003, ¶¶156-158, 169-173. For example, I pointed out that OCP teaches many advantages of using threads, including "support[ing] many different burst styles," "support[ing], concurrent activity and out-of-order completion of transfers," and "enabl[ing] busy targets...to prioritize requests." Ex. 1003, ¶157; Ex. 1006, 21-22, 46-47, 181. I also explained that threads allowed for the implementation of specific communication properties

such as ordering and priority. Ex. 1003, ¶169. Dr. Bagherzadeh does not dispute that a POSA would want these benefits or that these benefits were known.

45. Dr. Bagherzadeh argues that a POSA “would not be motivated to implement OCP’s thread-based teachings into Radulescu2004 because Radulescu2004 already teaches communication properties,” including ordering, and that a POSA “would not be motivated to use OCP’s threads in Radulescu2004 as a form of communication service identification because Radulescu2004 teaches that ‘the slave for which the transaction is destined is selected based on the address.’” Ex. 2070, ¶184. Dr. Bagherzadeh has failed to understand the purpose of the Radulescu2004 and OCP combination. Radulescu2004 teaches ordering and differentiated service classes within its NoC. Ex. 1005, 1-3. As indicated above, Radulescu2004 discloses “backward compatibility” with “existing on-chip communication protocols” such as “OCP.” Ex. 1005. A POSA would understand that OCP provides an example of how such threads are implemented in the presence of OCP. OCP masters use threads to indicate that a request has certain properties. For example, OCP teaches that its threads are used to indicate ordering, burst type and priority. Ex. 1006, 46, 181. Transfers within the same thread must “remain in order,” while those within different threads may be completed “out-of-order.” Ex. 1006, 46, 181. Threads may also be used by masters to communicate priority. Ex. 1006, 181.

46. As I discussed in my First Declaration, when the master uses a communication protocol that supports threading (such as OCP), a POSA would have been motivated to use threads as the communication service identification. Ex. 1003, ¶¶176-177, 164-173, 157-158. OCP teaches that threads have benefits, such as allowing for improved concurrent operation between a master and slave, that go beyond the ordering and differentiated service classes discussed by Radulescu2004. Ex. 1006, 21-22; Ex. 1003, ¶157. Threads can be used to assign properties to requests at a more granular level than address ranges. Using threads, the same master can send requests to the same memory region within a slave with different properties (e.g., ordering, burst type, priority).

3. OCP and Radulescu2004's Connections are Compatible

47. Dr. Bagherzadeh argues that a POSA would “understand that the ‘connection’ described in the ‘052 Patent and Radulescu2004 is different than the connection described by OCP” and that my opinions “rely on OCP’s connections being synonymous with the ‘052 Patent’s connections.” Ex. 2070, ¶¶187-188. I do not rely on OCP’s connections being the same as the ‘052 patent’s connections. The combination I discussed in my First Declaration does not call for the use of OCP’s connections in Radulescu2004’s system (as suggested by Patent Owner at Paper 38, 52). Ex. 1003, ¶¶155, 165, 174-177. I rely on Radulescu2004’s connections. *Id.* The combination does not call for the incorporation of OCP’s

interfaces into Radulescu2004's network or that OCP can natively interface with a NoC as Dr. Bagherzadeh suggests. Ex. 2070, ¶191. I point to OCP's teaching of a connection with an MConnID mapping to a thread with an MThreadID during the request phase only as yet more evidence that a POSA would have been motivated to map OCP's MThreadID to Radulescu2004's connections. Ex. 1003, ¶158. It shows that a POSA would have been motivated to provide mapping between connection level IDs and thread level IDs in interfaces. Ex. 1003, ¶158. Dr. Bagherzadeh argues that this mapping is backwards. Ex. 2070, ¶189. However, a POSA would understand that such mapping is an association of a MConnID and an MThreadID, and that the order would be irrelevant. In OCP, threads (identified by MThreadID) are communicated by Masters locally (Ex. 1006, 38-39, 46-47 ("The use of thread IDs allows two entities that are communicating over a OCP interface to assign transfers to particular threads.")), and connections (identified by MConnID) are used for end-to-end requests from masters to slaves (Ex. 1006, 47 ("Use connections whenever information about a request must be sent end-to-end from master to slave.")). Threads are mapped to connections because the various interfaces in an OCP path use different thread IDs. Ex. 1006, 47 ("thread IDs are re-assigned on each OCP interface in the path.").

48. Dr. Bagherzadeh argues that "OCP is a connectionless communication protocol." Ex. 2070, ¶187. Whether OCP is connectionless or whether it has

packets does not matter for the compatibility of OCP with Radulescu2004's system.

49. Dr. Bagherzadeh argues that “a POSITA would have understood that the disclosure in OCP that ‘any on-chip interconnect can be interfaced to OCP’ was incorrect at the time of the invention of the ‘052 patent because NoCs were not commercially available when OCP was first published.” Ex. 2070, ¶191. However, Radulescu2004 teaches that it is backward compatible with existing bus protocols including OCP. Ex. 1005, 1. For Radulescu2004's system to be backward compatible with OCP, a POSA would understand that Radulescu2004's network interfaces would be able to interface with OCP.

4. Any Deficiencies in Radulescu2004 Regarding Mapping Are Cured by OCP

50. Dr. Bagherzadeh argues that “OCP fails to cure the deficiencies of Radulescu2004 with respect to ‘mapping ... to a connection based on a set of connection properties’ because ... a POSITA would not combine OCP's teachings regarding ‘mapping’ with Radulescu2004.” Ex. 2070, ¶198. I disagree. As I discussed above, Radulescu2004 discloses the claimed mapping. But even under Dr. Bagherzadeh's narrow interpretation of the claim language, a POSA would have been motivated to make such a system in view of Radulescu2004 and OCP for the reasons I discussed in my First Declaration. For example, in my First Declaration, I explain OCP teaches MThreadID is mapped with a connection

identifier (MConnID), and that a POSA would have been motivated to have the OCP adapter in the network interface of Radulescu2004 map the MThreadID in requests to the analogous connections in Radulescu2004. Ex. 1003, ¶158. Mapping OCP threads to Radulescu2004's connections enables thread-based ordering and priority to be communicated through the NoC between system initiators and a system target and enables implementation of Radulescu2004's differentiated traffic classes. Ex. 1003, ¶177. A POSA would have been motivated to make such a mapping. The whole point of having connections with properties is to send requests for which particular properties are desired (e.g., ordering or high priority) over a connection with the desired property (e.g., ordering or high priority).

C. Ground 3: Anticipation by Weber

1. All Embodiments of Weber disclose a Single Integrated Circuit

51. As an initial matter, I note that Dr. Bagherzadeh agrees with me that Weber discloses sub-system 102 being on a single IC. Ex. 2070, ¶216 (“In FIG. 1 of Weber, the sub-system 102 ‘is a fully custom integrated circuit.’”). In my opinion, sub-system 102 alone is sufficient to anticipate claims 1-7. *See* Ex. 1003, Section X.C (e.g. ¶¶185, 190, 194-195, 199-201, 209-210, 219); *supra*, §VI.C.4.

52. Dr. Bagherzadeh disputes, however, that Weber discloses all of system 100 being on a single integrated circuit (IC). In support of his argument, Dr. Bagherzadeh references the embodiment of FIG. 1 discussed in the previous

paragraph, in which each sub-system 102, 104, 106, 108, and 110 is a distinct IC. Ex. 2070, ¶216. Weber, however, discloses that another embodiment of FIG. 1 integrates all of the sub-systems 102, 104, 106, 108, and 110 on a single IC. Ex. 1007, ¶0038 (“System 100 is an example of a system that *may consist of one* or more *integrated circuits* or chips.”).

53. Dr. Bagherzadeh argues that this sentence does not mean that sub-systems 102-110 can be implemented on one IC. Ex. 2070, ¶217. Instead, he argues that it merely refers to an embodiment in which “system 100 may only include one of the sub-systems (e.g., sub-system 102),” which would result in system 100 consisting of one IC. *Id.* Dr. Bagherzadeh’s interpretation, however, contorts that plain reading of the sentence which states that “**System 100** is an example of a system that *may consist of one or more integrated circuits* or chips” (Ex. 1007, ¶0038) rather than Dr. Bagherzadeh’s clearly incorrect interpretation that “**system 100** may *only* include *one of the sub-systems*.” Moreover, Weber separately states that each sub-system 102, 104, 106, 108, and 110 may be on a single IC, and also separately states that system 100 may include fewer sub-systems, so Dr. Bagherzadeh’s reading of this sentence makes it redundant with other statements in Weber. Ex. 1007, ¶0037 (“In one embodiment, *the sub-systems connected by communication bus 112 are separate integrated circuit chips*.”); (“Other embodiments may also *include fewer subsystems than the sub-systems*

shown in system 100.”). Also, my understanding of Weber is not based solely on this sentence. Weber contains other statements indicating that all sub-systems can be on a single IC. Ex. 1003, ¶185.

54. Dr. Bagherzadeh makes a series of unsupported statements about the integration of system 100 on a single IC and uses these statements to suggest that a POSA would not read Weber as disclosing all of the system 100 on a single IC. Dr. Bagherzadeh stated that such an integration:

- was “highly impractical and unusual”
- “would compromise either the performance of the CPU or the cost and density of DRAM”
- “would significantly increase the die size, rendering the chip too expensive and low performing,” and
- “would introduce thermal management issues.”

Ex. 2070, ¶¶219-224. He also stated that:

- “DRAM 106 is inherently incompatible with the logic processes of the integrated circuits 102 and 104 in FIG. 1 of Weber”
- “DRAM is typically upgraded more frequently than CPUs,” and
- if “all of system 100’s components were combined on a single integrated circuit or chip, the system would not operate efficiently because the system’s performance would suffer and it would be too unreliable.”

Id.

55. None of this suggests that Weber does not disclose all of the system 100 on a single IC. Integrating components together always involves tradeoffs, particularly with regard to cost, size, performance, and thermal management. The fact that these tradeoffs exist does not mean that Weber does not disclose

integrating the components together. There are probably many aspects of Weber that may be disadvantageous in certain situations, but that does not mean Weber does not disclose them. Moreover, it was known, even at the time of Weber, that integrating DRAM with a processor was both possible and beneficial. *See* Ex. 1053 (e.g. at abstract “A *microprocessor integrated with DRAM on the same die* has the potential to *improve system performance* by reducing the memory latency and improving the memory bandwidth.”)

56. Dr. Bagherzadeh also suggests that incorporating all the components is counter to Weber’s desire for a “low size or form factor,” because integrating the components would result in a chip that is “too large to meet the industry’s requirements.” Ex. 2070, ¶225. This is an unsupported claim. Dr. Bagherzadeh should know that every system in industry does not have the same requirement, and every system or class of systems would have their own constraints and requirements that a POSA would meet through integration. Weber clearly recognized that different systems have different requirements because it discloses that all or a subset of the components of system 100 may be integrated on a single integrated circuit. For example, contrary to Dr. Bagherzadeh’s statement, integrating all or a subset of the components of system 100 together on a single IC, as Weber discloses, *decreases* the overall size of the system as is well known by *reducing* the form factor of the entire system.

57. This was well-known in the industry. For example, an integrated circuit product from an industry leader in 2002 shows a single integrated circuit with DRAM, SRAM, ROM, FLASH, Microprocessor, DSP, Peripherals and ASIC. Ex. 1054 (Infineon Technologies, TriCore™ 1.3 32-bit Unified Processor Core, Architecture Overview Handbook, V1.3.3, May 2002). It states “Present and future trends for embedded systems include a *convergence of microcontroller and DSP architectures, as well as super-integration of memory and logic*. Embedded applications are evolving towards a *single System-On-a-Chip (SOC) comprising of a unified microcontroller-DSP core (32 bits), data and program memory (RAM, ROM, OTP, etc.), and custom application-specific logic (ASIC)*,” and “The *single core* provides virtual multiprocessing, *eliminating the need for multiple controllers and DSPs. On-chip memories enhance performance and reduce system power dissipation*, while the *integration of system peripherals and customer-specific logic increases the overall system performance at a reduced cost*.” Ex. 1054, 8.




Figure 1 System-on-a-Chip for Embedded Applications

Ex. 1054, 8.

58. Another example is a 1998 Texas Instruments product Technical Brief states (TMS320C6000 Technical Brief): “When *memory and peripherals are integrated with a CPU into one chip*, the *overall system cost is greatly reduced*, and *circuit board space is reduced*.” Ex. 1055, 12. This was also well-known even as far back as 1986. Ex. 1056, 7 (“The towed array project involves the design of a 3500-gate CMOS array for use in the second-generation underwater data acquisition module. The *gate array replaces 10 medium- and large-scale integrated circuits* used in the initial design, *reduces the size and power consumption of the circuit*, and permits the addition of new functional capabilities.”). Yet another example is a publication in 2001 entitled “Platform Tuning for Embedded Systems Design” published in Computer Magazine states that “*System-on-chip technology can now integrate* entire digital systems that *used to*

occupy one or more boards—systems consisting of *numerous components like microprocessors, memories, co-processors, and peripherals—onto a single chip.*”

Ex. 1057, 1. Yet another example is M. Slater, "The microprocessor today," in *IEEE Micro*, vol. 16, no. 6, pp. 32-44, Dec. 1996, stating: “Instead, embedded application designers would like to *reduce system costs by integrating more functions on the same chip with the microprocessor*. The logical end point of this evolution is a *complete system on a chip*. Technology is *reaching a point where chips can integrate even significant amounts of memory*. For example, eliminating half the DRAM array from a 64-Mbit DRAM still leaves 4 Mbytes of memory and room for millions of logic transistors... As *embedded microprocessors evolve* toward systems on a chip, they *inevitably become more specialized. Different applications have different needs for memory, peripheral controllers, and interfaces* to the external world. The *desire for highly integrated system chips* is increasing the demand for building block microprocessors that *can function as parts of application-specific integrated circuits (ASICs)*.” Ex. 1058, 11-12.

59. Dr. Bagherzadeh cites no references to support his suggestion that such integration would be “too large,” or what “industry[] requirements” he is referring to. On the other hand, Weber fully recognizes that different industry applications would have different constraints and optimization criteria and as such

disclosed embodiments with one or more of the subsystems of system 100 integrated on a single integrated circuit.

60. Dr. Bagherzadeh also ignores the fact that Weber discloses “[o]ther embodiments may also include fewer subsystems than the sub-systems shown in system 100.” Ex. 1007, ¶37. Based on this and the overall teachings of Weber, a POSA would at once envisage an embodiment of Weber including only, for example, sub-systems 102 and sub-system 104 or sub-system 102 and FPGA 110 on a single IC. Such an embodiment is yet another way in which Weber meets the limitations of claims 1-7.

2. Weber has Same Structure as ‘052 for Mapping Means

61. In my first declaration, I asserted that the adapter unit is the item in the ‘052 that corresponds to the mapping means. Ex. 1003, ¶85. Dr. Bagherzadeh, however, asserted that both a sequentializing unit and a desequentializing unit are required. Ex. 2070, ¶231. If disclosure of a sequentializing unit and/or desequentializing unit is required, Weber discloses the same amount of structure as the ‘052. The ‘052 illustrates black boxes for both of these units and then proceeds to describe the functions performed by those black boxes. Ex. 1001, FIGs. 2A, 2B; 6:48-55 (“The *basic operation of such a shell associated to a master as shown in FIG. 2A is to sequentialize* commands cmd and their flags flag, addresses adr, and write data wr_data in request messages msg, and *to desequentialize* messages msg

into read data rd_data, and write responses wr_resp. The basic operation of a shell associated to a slave *as shown in FIG. 2B* is reversed to that of the shell according to FIG. 2B.”). Weber discloses the same thing. Weber describes embodiments that perform the function of multiplexing, and a POSA would understand that such multiplexing happens in the interface module. Ex. 1007, ¶¶0036, 0038, 0049, 0050. Additionally, as discussed below, Weber’s disclosure of multiplexing indicates to a POSA that demultiplexing also occurs in the interface module.

62. Dr. Bagherzadeh asserts that “A POSITA would understand that Weber does not disclose a desequentializing unit because once signals are multiplexed, there is no reason to de-multiplex them.” Ex. 2070, ¶231. I disagree. On the contrary, a POSA would understand that in a bidirectional system where signals are transmitted via multiplexing information (i.e. sequentializing) using the same modules and shared bus, the reception of information would require demultiplexing (i.e. desequentializing). Weber discloses that: “Shared *communications bus 1012* includes separate address, data, and control wires. Other embodiments *may include multiplexed address, data, and control signals that share a wire or wires*. Such an embodiment would provide high per-wire transfer efficiency because *wires are shared among address and data transfers*.” Ex. 1007, ¶50. As shown in FIG. 4 of Weber’s discloses that the information (e.g. address, data, and control) are multiplexed onto shared bus 1012 and thus a POSA

would understand that a multiplexing (i.e. sequentializing) is required. A POSA would also understand that 1012 is a bidirectional bus and, for example, initiator interface module 1004, can receive information that has been already multiplexed on shared bus 1012 and thus, a POSA would understand that de-multiplexing (i.e. de-sequentializing) is disclosed by Weber. As shown in FIG.4 and FIG. 11 and their descriptions, a POSA would understand that Weber discloses multiplexing on the shared bus (e.g., bus 1012 of FIG. 4), and parallel (non-multiplexed) signals on the interconnects 1010. Ex. 1007, ¶¶0050, 0074-0082 (“Shared *communications bus 1012 ... may include multiplexed address, data, and control signals that share a wire or wires*” and “FIG. 11 shows an *interconnect 1010*, which is a *point-to-point interconnect as shown in FIG. 4*.... Referring to FIG. 5, a *master entity may be, for example, initiator functional block 1002 or target interface module 1006*. A *slave entity may be, for example, initiator interface module 1004 or target functional block 1008*,” and “*Response (Resp) and DataIn are signals sent from slave 1104 to master 1102*. Resp indicates that a transfer request that was received by slave 1104 has been serviced,” and “*Signals Clock, Cmd, Addr, DataOut, ReqAccept, Resp, DataIn, and RespAccept*, in one embodiment, make up a basic set of interface module signals.”). For example, as shown in FIG. 4 and FIG. 11, a POSA would understand that this would require the initiator interface module 1004 to multiplex the parallel signals on interconnect 1010 from initiator

functional block 1002 for placement onto the bus 1012 and would also require target interface module 1006 to de-multiplex the signals from bus 1012 for placement as parallel signals onto the interconnect 1010 to the target functional block 1008. Ex. 1001, FIGs. 4, 11 (annotated below).




FIG. 4




FIG. 11

63. Dr. Bagherzadeh also argued that Weber does not disclose sequentializing in an integrated circuit, based on Weber’s “more appropriate for chip-to-chip communications” statements. Ex. 2070, ¶232. I addressed this in my first declaration, when I indicated that such statements would be understood by a POSA to communicate that such multiplexing is performed *both on and off* an IC. Ex. 1003, ¶225. Dr. Bagherzadeh tries to narrow the language of Weber by

suggesting Weber states multiplexing *is* (always) more appropriate for chip-to-chip communications. Ex. 2070, ¶232. What Weber actually states is that there is a tradeoff between “higher performance” and “extra wires,” and this tradeoff *may* cause multiplexing to be more appropriate for chip-to-chip communication. Ex. 1007, ¶0036; *see also, id.*, ¶0050. For example, a POSA would understand Weber to disclose using multiplexing on an IC when “high per-wire transfer efficiency” is desirable on the IC. *Id.*

64. Dr. Bagherzadeh argues that I opined in my first declaration that such multiplexing would occur on the buses 112, 114 and not in the “mapping means (A).” Ex. 2070, ¶233. What I actually state is that the signals on the buses 112, 114 are “multiplexed,” but that the “multiplexing” is performed by the interface modules for those buses 112, 114. Ex. 1003, ¶¶224-225. I clearly state that “The interface modules of Weber perform the function of the claimed mapping,” “a POSA would at once envisage that the interface modules of system 100 were intended to perform such multiplexing,” and “the interface modules of system 100 are the ‘mapping means’.” *Id.* Dr. Bagherzadeh does not appear to dispute that the interface modules of Weber would perform the multiplexing.

3. Weber Discloses an Address Range

65. Dr. Bagherzadeh criticizes my reliance on FIG. 3 of Weber, because FIG. 3 discloses a target interface module. Ex. 2070, ¶239. This criticism appears

to be based on Dr. Bagherzadeh's mistaken opinion that the address range must be sent from the initiator. As discussed above in Ground 1, the claims do not require the address range to be sent from the initiator. *Supra*, §VI.A. The claims merely require that the first module "requests at least one communication service ... **based on** ... at least one address range." *Id.* FIG. 3 discloses the request being "based on" an address range as discussed in my first declaration. Ex. 1003, ¶193.

66. Next, Dr. Bagherzadeh criticizes my use of the word "capability." Ex. 2070, ¶¶240-241. I did not intend to suggest that Weber does not disclose actually sending a request that meets the claim language. I clearly state: "Weber describes that the function blocks provide an 'address (Addr) signal' with the request" and that "Weber's request includes an address signal" (among other instances) and then go on to explain how sending such a request meets the claim language. Ex. 1003, ¶193.

67. Dr. Bagherzadeh then argues that I "fail[] to provide any explanation as to why a POSITA would understand the 'client' in 'client address base' refers to the target." Ex. 2070, ¶243. A POSA would understand that "client" as used in this instance refers to the client of the target interface module, which is the target module. Thus, client address base in FIG. 3 refers to the target module.

68. In Weber, "FIG. 3 is a block diagram of an embodiment of a **target** interface module 900," (Ex.1007, ¶0046), while "FIG. 2 is a block diagram of an

embodiment of an *initiator* interface module 800,” (Ex.1007, ¶0043). Dr. Bagherzadeh attempts to intentionally confuse the issue by jumping to the use of “client address base” in FIG. 2. Ex 2070, ¶243. FIG. 2 describes an example of an *initiator* (rather than, as discussed above, *target*) interface module. Ex. 1007, ¶0043. Analogous to what I described above, the use of “client” in FIG. 2 refers to the client of the *initiator* interface module, which is, of course, the initiator module. Therefore, the use of “client” in FIG. 2 does in fact refer to the initiator, but the use of “client” in FIG. 3 refers to the target and not the initiator as Dr. Bagherzadeh incorrectly asserts.

69. Next, Dr. Bagherzadeh argues that I fail to provide any explanation as to why a POSITA would understand that the discussion of the address base or the external address portion indicates that the target has a range of addresses. Ex. 2070, ¶¶244-246. As I have explained in my first declaration (*see* Ex. 1003, ¶193), by disclosing that there is an external address portion, that indicates that there is also an internal address portion. This communicates to a POSA that there is a range of addresses within the target that can be addressed. Similarly, disclosing that there is a “base” address, communicates to a POSA that there is a range of addresses within the target that can be addressed.

70. Dr. Bagherzadeh also makes several arguments that “a POSITA would understand that the address signal (“Addr”) does not identify the target itself.” Ex.

2070, ¶¶249-251. Dr. Bagherzadeh, however, is confusing the issues. The claims do not require that the address “identify the target.” The claims require that the address range is “for identifying one or more second processing modules.” The address range is for identifying the target as discussed in my first declaration. Ex. 1003, ¶193. In any case, a POSA would also understand the address itself to be “for identifying one or more second processing modules” for the reasons discussed in my first declaration. *Id.*

4. Weber Discloses an Interconnect Means Under All Constructions

71. Dr. Bagherzadeh does not dispute that Weber discloses an interconnect means if the corresponding structure for an interconnect means is a direct bus as I opined in my first declaration. Ex. 1003, ¶82. Weber also meets Dr. Bagherzadeh’s construction “multi-hop interconnection between modules separated by one or more network nodes.” Ex. 1003, ¶¶199-205. Each of the interface modules 120, 122, 123, 124, 125, 126, 127, 128, and bridge F in system 100, as well as each of communication subsystems 1435, 1440 of FIG. 14, creates a “hop” and is a “network node.” *Id.*; Ex. 2068, 22:8-12, 26:6-12, 111:20-114:13. Thus, sub-system 102, system 100, and FIG. 14 disclosed in Weber each provide multiple hops from a master to a target, which are also separated by one or more network nodes.

72. For example, sub-system 102 includes a “multi-hop interconnection between modules separated by one or more network nodes,” because the functional

blocks A, C, E, and F are “separated by one or more network nodes” due to the respective interface 120, 122, 123, 124 in front of each of those blocks. The interconnect coupling those blocks together is also “multi-hop” because the communications must hop to and from these interfaces 120, 122, 123, 124 to get to and from those functional blocks. *See* Ex. 1003, ¶194-197, 199-201; Ex. 2068, 22:8-12, 26:6-12, 111:20-114:13. For example, I showed an example of this at ¶194 of my First Declaration, for the path from A to C:




FIG. 1

This is multi-hop because the communication hops from Functional Block A to Interface 120, then hops again from Interface 120 to Interface 122, and hops again from Interface 122 to Functional Block C, and Interfaces 120 and 122 are network nodes. Ex. 1007, ¶¶0040, 0048, FIG. 1. Figure 4 also shows a generalized example of this, as I also showed at ¶195 of my First Declaration. Similarly, system 100 overall is another disclosure of a “multi-hop interconnection between modules

separated by one or more network nodes” because it shows, for example, a path from block A of sub-system 102 to block A of sub-system 110 that hops to/from interface 120, interface 124, bridge F, and interface 128 between block A of sub-system 102 and block A of sub-system 110. Ex. 1003, ¶¶199, 201; Ex. 1007, ¶¶0040, 0048, FIG. 1. I showed this in ¶199 of my First Declaration:




FIG. 1

Figure 14 is yet another disclosure of a “multi-hop interconnection between modules separated by one or more network nodes” because there is a hop from the initiator function block 1410 across the first communications sub-system 1435 and a hop across the second communications sub-system 1440 to the target functional block 1430. Ex. 1003, ¶¶203-205.

73. Dr. Bagherzadeh disputes that Weber discloses an interconnect means under his construction of the term. First, Dr. Bagherzadeh appears to suggest that Weber discloses *only* a bus. Ex. 2070, ¶¶256-257. As I discussed in my first

declaration, Weber discloses more than just a single bus. Weber discloses an interconnect including multiple shared-buses and point-to-point interconnects, connected together with interfaces and a bridge. Ex. 1003, ¶¶199-205.

74. Dr. Bagherzadeh next makes several arguments about limitations that are not in the claims. He argues that Weber does not disclose “protocol conversion” that “each routing path is directly connected,” that it does “not send packets,” that it does not “choose the channel for the data flow,” and that it does not disclose any “routers or switches.” Ex. 2070, ¶¶258-274. Dr. Bagherzadeh’s explanations for many of these limitations are illogical and therefore in many cases it is unclear what exactly he contends is required. In any case, I do not believe the claims require any of these aspects.

75. With respect to packets, none of the claims use the word “packet” and Dr. Bagherzadeh could not identify any language in the specification that indicated that the claims required packets. Ex. 1050, 65:24-74:12.

24 Q. You'll agree with me that the word "packets"

25 don't appear in the claims of the '052 patent?

1 A. I don't have to check. It doesn't.

.

.

.

6 Q. Okay. You will agree with me that packets are

7 not referred to in the '052 patent outside of the

8 section entitled "Description of the Preferred

9 Embodiments"?

10 A. I think -- I think that's what it shows from my

- 11 search. It's those columns that I stated earlier. And
12 we established that -- that's about it.

The IEEE dictionary also suggests that packets are not required for all networks.

It's definitions of "network" do not require packets and it indicates that the more specific term "packet switching network" is used for networks that use packets.

Supra, ¶80; Ex. 1052, 7 ("**packet switching network** A network that uses packet switching techniques for transmission of data. *Synonyms*: packet switched network, packed data network.")

76. Even if packets were required, Weber discloses packets. Neither the '052 patent, nor Dr. Bagherzadeh provides any particular requirement that a communication must meet in order to be a "packet." In general, "packet" encompasses any collection of information sent in a single operation. *See* Ex. 1052, 787-788; Ex. 1051, 358. Weber discloses sending a collection of information together in a single operation and therefore discloses packets. Ex. 1007, ¶¶0049-0054. To the extent packet requires something more specific, like sending information sequentially including an address and data, Weber's disclosure of multiplexing, which places the address, data, and control signals on the same wire or wires, is a disclosure of packets. Ex. 1007, ¶¶0036, 0050.

77. With respect to protocol conversion and network node, none of the claims recite protocol conversion and Dr. Bagherzadeh could not identify any language in the specification that defined "network node." Ex. 1050, 60:22-61:3.

22 THE WITNESS: I -- I don't see a sentence
23 saying, Now we're defining network nodes, no.
24 If you look at the body of this and the
25 figures, you will understand what the network node is.
1 It's very clear to me in that respect. But there is not
2 a single sentence saying, Now I'm defining what a
3 network node is.

Even if protocol conversion were required, Weber discloses protocol conversion.

Neither the '052 patent nor Dr. Bagherzadeh provides any particular requirement that must be met for protocol conversion to occur. Weber, however, indicates that interconnect 115 and buses 112, 114 use different protocols. Ex. 1007, ¶¶0049, 0074 (“Interconnect 1010 includes additional signals as compared to the protocol described with reference to FIG. 5”); Ex. 1003, ¶¶194-197, 211. Additionally, Weber provides further disclosure in other embodiments that interconnects 115, 112 and 114 are different protocols. As I have discussed above in §VI.C.2, a disclosed embodiment is point-to-point buses 115 and 114 using parallel signals while shared bus 112 is shared signals. Thus, protocol conversion occurs at interface modules 120, 122, 123 and 124 between respective interconnects 115 and bus 114 and at bridge F between interconnect 115 and bus 112. Accordingly, even if protocol conversion were required to be a “network node,” Weber’s interface modules 120, 122, 123, 124, and bridge F are “network nodes.”

78. With respect to path variability and the ability to choose the channel, again none of the claims recite this and Dr. Bagherzadeh admitted that multiple paths were not required. Ex. 1050, 202:2-203:15:

2 Q. Okay. Well -- and maybe your question's --

3 that suggests a more basic question.

4 Does a **multi-hop network mean that you have to**

5 **have more than one path** to get from a master to a slave?

6 A. **Not required.** You can if -- based on the

7 design of the NoC, you have that option because you have

8 these short links to other routers. And that's the

9 beauty of NoC compared to a bus or, in this case, the

10 way that he handles the communication between entities.

11 So you have that option. If you don't want to

12 use it, you don't use it. That's the routing algorithm

13 that we worked on, basically. You can have a very

14 deterministic routing, or you can have an adaptive

15 routing. Weber doesn't have that because it's not based

16 on that.

17 Q. I think we talked about this earlier; and I

18 believe you indicated that if I had a very simple system

19 that had two modules and a single router, that that

20 could still practice the method of, say, Claim 6 of the

21 '052 patent?

22 MR. CHRISTOPHER: Object to form.

23 THE WITNESS: **One or more nodes.** It's in the

24 description. And I had explained to you that the link

25 to the router is one hop; the other one is the other

1 one.

2 BY MR. VANDENBURGH:

3 Q. Okay.

4 A. It -- it -- it **clearly says one or more.**

5 Q. But you agree with me in that circumstance

6 there would be **only one path** for the data to take from

7 the initiator module to the target module?

8 A. **Yes.** But comparing that to a bus I think to me

9 misses the point that the scalability of NoC --

10 everybody talks about scalability. If you want improve

11 the scalability with one router, then it's really not a
12 fair game. Because bus is limited. But the NoC can
13 expand. Most of the NoCs I've seen have -- but the hops
14 count matches the one node. That's the most -- you
15 don't an OC11 router, usually. But hops count.

The common understanding of a “network” also does not require path variability or the ability to choose the channel. *See* Ex. 1052, 5-6; Ex. 1051, 10. Even if path variability and the ability to choose the channel were required, Weber discloses it. With respect to FIG. 14, Weber discloses that there can be variability in paths between an initiator and a target and that data is directed through those variable paths to the intended target. Ex. 1007, ¶¶0084, 0100, 0107-0115, FIGs. 15B, 18.

79. With respect to a router or a switch, again, none of the claims use the words “router” or “switch” and the common understanding of “hop” does not require a router or switch. *Supra*, §V.B. Dr. Bagherzadeh’s position that a hop is only created by a router or switch is unsupported. Ex. 2070, ¶¶114, 123, 254. None of the references he cites indicate that a router or switch is the *only* thing that creates a hop. Ex. 2070, ¶¶144 (citing 2005), 122 (citing Ex. 2073).


80. Routers and switches are not required for all networks either. For example, the IEEE Dictionary of 2000 provides nine definitions of “network” none of which even mention switches or routers and most of which clearly indicate that the components that make up the network are not limited. Ex. 1052, 5-6. The definitions include:

- (1) A combination of elements or devices
- (2)(A)(data transmission) A series of points interconnected by communication channels
- ...
- (6) A collection of interconnected hosts
- (7) A network is any set of devices or subsystems connected by links joining (directly or indirectly) a set of terminal nodes
- (8) An arrangement of components, or nodes, and interconnecting branches.


Id.

81. Webster’s Telecom Dictionary of 2008 similarly defines network as “A fabric of elements that work together much as the fabric or mesh of a net to support the transfer of information.” Ex. 1051, 10. It goes on to specify that “A network includes all links, amplifiers, and repeaters, multiplexers, switches, routers and other intermediate devices involved in establishing maintaining and termination a session between a transmitter and a receiver.” *Id.*


82. The patent also does not suggest a “network” must include a router or switch. The patent states that a network “may” include routers and “may” include switches. Ex. 1001, 4:48-50, 5:24-25. FIG. 4 of the patent also illustrates a network and does not show any switches or routers, only network interfaces. Ex. 1001, FIG. 4:




A POSA would understand that the network N can be implemented without using switches or routers, for example, by using one or more buses interconnected with bridges. *See* Ex. 1051, 3-4 (“The network was in the form of a physical bus topology, as illustrated in Figure B-5. All devices connected to the cable, and communicated over a single, shared channel on a shared electrical circuit. ... Ethernet segments could connect through bridges”) (Figure B-5 Reproduced Below).



83. Dr. Bagherzadeh also opines that “No POSITA would call [the implementation on the left of the Figure in Ex. 2073 at 15 a “multi-hop interconnect.” Ex. 2070, ¶265; Ex. 2073, 15 (Figure Reproduced Below).



However, the Data & Computer Communications textbook submitted by the Patent Owner describes such a configuration as having multiple hops. Ex. 2010, 460-461 (Figure Reproduced Below).



The textbook describes:

Figure 13.16 also shows that there are two routes between LAN A and LAN E ... if station 1 transmits a frame on LAN A intended for station 5 on LAN E, then either bridge 101 or bridge 107 could forward the frame... if frame travels through bridge 101, it must suffer two hops.

Ex. 2010, 461 (emphasis added). The textbook continues: “typically the route with the least number of hops is selected.” *Id.* (emphasis added). Dr. Bagherzadeh is incorrect that a POSITA would not call such an implementation “multi-hop.”

84. Even if a router or switch were required, Weber discloses it. As I stated in my first declaration, a POSA would envisage that communication sub-

systems 1435 and 1440 in Weber are routers. Ex 1003, ¶204. Dr. Bagherzadeh does not dispute this. Weber, for example, discloses that data received at interface 1450 of communications sub-system 1435 is directed through interfaces 1460 and 1480 if target 1430 is the destination. Ex. 1007, ¶0097; Ex. 1003, ¶¶199, 203-204. Thus, communication sub-system 1435 routes data and is a router. The fact that communication sub-systems are “loaded with a mapping” does not change this.

5. Weber’s Bus-with-Interfaces is a “Network”

85. As mentioned above, Weber’s system 100 discloses an interconnect including multiple links/buses (112, 114, 115), multiple interfaces (120, 122, 123, 124, 125, 126, 127, 128), and a bridge (F) interconnecting a plurality of functional blocks. This bus-with-interface topology is a “network” as that term is used in the ‘052 patent. As described above, each of the interfaces (and the bridge F) create a hop making Weber a “multi-hop interconnect.” *Supra*, §V.B, §VI.C.4. Each of the interfaces (and the bridge F) are also part of the network and therefore “network nodes,” thus the functional blocks are “separated by one or more network nodes.” *Supra*, §V.B, §VI.C.4; Ex. 1003, ¶¶193-196, 209.

86. The network of ‘052 patent includes interfaces that are network nodes (that introduce a hop) that perform functionality that “changes the communication” and “separate[]” the communication modules (e.g. Master, Slave) from each other rather than having “direct interconnects” between them. The ‘052 patent discloses:

“Introducing *networks* as on-chip interconnects radically *changes the communication* when compared to *direct interconnects*, such as buses or switches. This is because of the *multi-hop nature of a network*, where communication modules are *not directly connected*, but *separated by one* or more *network nodes*.” Ex. 1001, 2:32-37. For example, it is the interfaces that “decouple computation from communication,” enable “re-use of applications and architectures,” and enable combination of modules using different protocols, all major described benefits the ‘052 patent ascribes to a “network” (“... NoCs also *decouple computation from communication*, ... NoCs *achieve this decoupling* because they are traditionally designed using protocol stacks, which provide *well-defined interfaces separating communication service usage from service implementation*,” (Ex. 1001, 1:42-47, *see also* 2:13-24, 2:43-53, 2:63-3:2). The inclusion of the ‘052 interfaces means the modules are not “directly connected” and instead communicate “remotely via network nodes” and are “separated by one or more network nodes.” *Id.*, 1:50-54; 2:33-38. Interfaces also enable multiple buses to be interconnected together, which reduces the communication bottleneck on a single bus, changes the arbitration from centralized to distributed, and makes the interconnect more scalable. *Id.*, 1:29-32, 40-42, 54-58, 2:38-42. The interfaces can also handle out-of-order transactions and end-to-end flow control. *Id.*, 1:56-57. Examples of functionality in the ‘052 interfaces to perform this “separate[ion]”

include: “The network *interfaces NI* are responsible for *packetization/depacketization*, for *implementing the connections and services*, and for *offering a standard interface (e.g., AXI or OCP)* to the IP modules connected to the NoC. In particular, the network *interfaces NI* may implement *adapter units A* to existing on-chip protocols, such as AXI, OCP and DTL, to seamlessly connect existing IP modules to the NoC. These adapter units A may also be arranged between the IP blocks and their associated network interfaces,” (Ex. 1001, 6:30-38, *see also* 6:56-7:14, 7:28-39), “...mapping means is *arranged in said at least one network interface*, such that the processing modules can continue with their dedicated operations *without having to deal with the mapping of services*,” (Ex. 1001, 3:30-34, *see also* 4:4-13, 8:7-36), “The communication *between* said at least one first of said plurality of processing modules and said interconnect means is *controlled* by at least *one network interface* associated to said at least one first of said processing modules,” (Ex. 1001, 3:54-57), “FIG. 1 shows a network on chip according to a first embodiment. Between *network interfaces and its associated master M and slave S* an adapter unit A is arranged, respectively, wherein the *adapter units A* may also be arranged in the network interface. ... The *network interfaces NI* are provided to manage the communication between the respective modules M, S and the network N, so that the modules can perform their dedicated operation without having to deal with

the communication with the network or other modules.” (Ex. 1001, 5:6-23),
“...In other words, *the adapter units A translate the requested communication service into messages msg, which are processed by the network interface NI and are translated into flits fit to be transmitted over the network N, in particular the adaptor unit A maps the identification means of these protocols to connections.* ...,” (Ex. 1001, 5:53-6:3). Indeed, the Wingard patent (U.S. Pat. No. 5,948,089) (incorporated by reference in Weber) teaches that interfaces in a bus system provide the key attributes of a “network” identified in the ‘052 patent, including decoupling computation and communication, enabling re-use of applications and architectures, and enabling combination of modules using different protocols. Ex. 1008, 4:20-26, 15:43-59.

87. This all suggests that the “direct” interconnects referenced in the ‘052 patent are interconnects without interfaces, e.g., a point-to-point wires or a bus without interfaces, and conversely that a bus with interfaces is not direct and is therefore a network as used in the ‘052 patent. Dr. Bagherzadeh acknowledges that Weber discloses an interconnect that “provides the flexibility missing from traditional buses.” Ex. 2070, ¶296. This is further evidence that Weber discloses a “network” as that term is used in the ‘052 patent.

6. Weber's Interfaces are "Network" Interfaces

88. Weber's interfaces are network interfaces for the simple reason that Weber discloses a network as discussed above and the interfaces are for that network. *Supra*, §VI.C.4, VI.C.5, Ex. 1003, ¶¶193-196, 209. Dr. Bagherzadeh raises the "protocol conversion" argument again with respect to the interfaces of Weber. Ex. 2070, ¶¶277-282. Again, however, Dr. Bagherzadeh does not explain why the recitation in the claims requires protocol conversion and, even if it did, Weber discloses protocol conversion by its network interfaces as discussed above. *Id.*

D. Ground 4: Weber and Radulescu2004

1. It Would Have Been Obvious to Implement Weber-Radulescu2004 Combination on a Single Integrated Circuit

89. In my first declaration, I opined that a POSA would have been motivated, with a reasonable expectation of success, to combine Weber and Radulescu2004, in a way that meets all of the claimed limitations. This includes the limitation for an "integrated circuit." It would have been obvious to implement the combination on an integrated circuit for the reasons discussed above and in my first declaration with respect to Weber's disclosure of an integrated circuit. *Supra*, §VI.C.1; Ex. 1003, ¶¶185, 191, 233-234. Dr. Bagherzadeh disagrees that putting them all on an integrated circuit would have been obvious but provides no further reasoning other than what he provided in Ground 3.

2. The Reasons Provided in My First Declaration Do Provide Motivation to Combine Embodiments of Weber

90. Dr. Bagherzadeh asserts that because Weber touts or otherwise provides benefits over traditional buses a POSITA would not have been motivated to combine Radulescu2004 with Weber, because Weber already provides these benefits. Ex. 2070, ¶¶293-299. The fact that Weber provides these benefits over traditional buses, however, does not mean that a POSA would not be motivated to make further improvements to scalability, performance, concurrent data transmissions, fault tolerance, and power efficiency. In fact, Weber's statements relating to the desirability of improved scalability and performance in an interconnect, support the position that these are things that motivate a POSA to make changes to an interconnect.

91. Dr. Bagherzadeh also suggests that the potential disadvantages of Radulescu2004's NoC need to be considered. Ex. 2070, ¶¶301-304. I did take into account any potential disadvantages because I expressly consider Radulescu2004's NoC as compared to a bus and conclude that there can be (on balance) advantages for Radulescu2004's NoC. Ex. 1003, ¶¶238-241. In any case, the mere presence of tradeoffs alone would not discourage a POSA from trying the proposed combination.

92. One of PO's own references provides further evidence of the known interchangeability of Radulescu2004's NoC with the bus topology of Weber. Ex.

1021, titled “A generic architecture for on-chip packet-switched interconnections”

states that switching networks are an alternative to a shared bus:

We explain why the shared bus, which is today’s dominant template, will not meet the performance requirements of tomorrow’s systems. We present an alternative interconnection in the form of switching networks.

Ex. 1021, Abstract; *see also, Id.*, §7 (including Table 3 reproduced below).

Bus Pros & Cons		Network Pros & Cons	
Every unit attached adds parasitic capacitance, therefore electrical performance degrades with growth.	-	+	Only point-to-point one-way wires are used, for all network sizes.
Bus timing is difficult in a deep submicron process.	-	+	Network wires can be pipelined because the network protocol is globally asynchronous.
Bus testability is problematic and slow.	-	+	Dedicated BIST is fast and complete.
Bus arbiter delay grows with the number of masters. The arbiter is also instance-specific.	-	+	Routing decisions are distributed and the same router is reinstanciated, for all network sizes.
Bandwidth is limited and shared by all units attached.	-	+	Aggregated bandwidth scales with the network size.
Bus latency is zero once arbiter has granted control.	+	-	Internal network contention causes a small latency.
The silicon cost of a bus is near zero.	+	-	The network has a significant silicon area.
Any bus is almost directly compatible with most available IPs, including software running on CPUs.	+	-	Bus-oriented IPs need smart wrappers. Software needs clean synchronization in multiprocessor systems.
The concepts are simple and well understood.	+	-	System designers need <i>reeducation</i> for new concepts.

Table 3: The bus-versus-network arguments

A POSA would have understood that Radulescu2004’s NoC could be used in place of Weber’s bus topology and would have known how to implement such a modification.

3. Sufficient Detail was Provided for a POSA to Combine Weber and Radulescu2004

93. Dr. Bagherzadeh also alleges that I failed to “specifically identify the resulting combination of Weber and Radulescu2004,” but then he admits that I provided several “permutations” and he fails to identify anything that is missing.

Ex. 2070, ¶305. A POSA would understand how to combine Weber and Radulescu2004 as I described in the first declaration. Ex. 1003, ¶¶243-251.

94. Dr. Bagherzadeh then asserts that a POSA would not be motivated to use Radulescu2004's NoC with Weber because Weber's interfaces communicate in a different protocol than Radulescu2004's network. Ex. 2070, ¶¶306-310. Both Weber and Radulescu2004, however, teach that the interfaces should communicate in the protocol of the interconnect. *Supra*, §VI.C.1; Ex. 1005, §4. Based on these teachings and the general understanding that interfaces should be able to communicate on the mediums to which they are interfacing, a POSA would understand that the interfaces of the resulting combination should communicate in the same protocol as the NoC and would know how to do so.

95. Dr. Bagherzadeh then makes yet another argument based on the mistaken conclusion that the address range must be present in the request from the initiator. Ex. 2070, ¶¶311-312. The claims contain no such requirement as discussed above. *Supra*, §VI.C.1. Thus, the target address range discussed in the '089 patent (incorporated by reference in Weber) is sufficient to meet the claim language. Ex. 1003, ¶¶246-248.

96. Dr. Bagherzadeh then argues that the resulting combination fails because Radulescu2004 independently fails to disclose certain elements of the claim. Ex. 2070, ¶¶316-318, 322. Dr. Bagherzadeh does not appear to dispute that

Weber teaches including these items, and that alone is sufficient motivation to implement the elements in the combination. With respect to Radulescu2004, I have already addressed Radulescu2004's teachings related to these elements in Ground 1, so I will not address them again. *Supra*, §VI.A.

97. In paragraph 320, Dr. Bagherzadeh suggests that both Radulescu2004 and Weber teach identifying a particular slave without using an address range. Ex. 2070, ¶320. I disagree for the reasons provided above in Grounds 1 and 3. *Supra*, §VI.A.1; §VI.C.3.

98. Dr. Bagherzadeh similarly asserts that Radulescu2004 also “allow differential traffic classes through the connid signals” without the use of Thread IDs. Ex. 2070, ¶321. Dr. Bagherzadeh does not appear to dispute that Weber teaches including these items, and that alone is sufficient motivation to implement the elements in the combination. Furthermore, a POSA would have been motivated to use Thread IDs in Radulescu2004. *Supra*, §VI.B.2.


E. IEEE

99. I note that Radulescu2004 was published in the proceedings of an IEEE-sponsored conference (the Design, Automation and Test in Europe Conference and Exhibition (Date'04)). Ex. 1005, 1. To a POSA in early 2004 (and still today), IEEE-sponsored conferences were one of the top sources of information regarding developments in relevant technologies, including SoC

design—even more so than journal articles. A POSA in this timeframe (and still today) would have frequently browsed IEEE conference paper publications online, and also attended IEEE conferences, for information on relevant technologies including SoC design.

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

Dated: January 17, 2025

A handwritten signature in blue ink, appearing to read "Marwan Hassoun", is written over a horizontal line.

Dr. Marwan Hassoun