

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

QUALCOMM INCORPORATED
Petitioner

v.

NETWORK SYSTEM TECHNOLOGIES, LLC
Patent Owner

Patent No. 7,594,052

Inter Partes Review No. IPR2024-00343

DECLARATION OF DR. NADER BAGHERZADEH

TABLE OF CONTENTS

I. INTRODUCTION.....	1
II. EXPERT QUALIFICATIONS AND BACKGROUND	2
III. MATERIALS CONSIDERED	11
IV. UNDERSTANDING OF LEGAL STANDARDS.....	13
A. Prior Art.....	13
B. Anticipation	14
C. Obviousness.....	16
D. Claim Construction.....	20
V. PERSON OF ORDINARY SKILL IN THE ART.....	22
VI. TECHNOLOGY OVERVIEW	24
A. Technical Background.....	24
B. Overview of the ‘052 Patent.....	29
VII.OVERVIEW OF PETITIONER’S REFERENCES	40
A. Rădulescu2004 (Ex. 1005).....	40
B. OCP (Ex. 1006)	42
C. Weber (Ex. 1007)	44
VIII. CLAIM CONSTRUCTION	46
A. “interconnect means” (claims 1, 6, and 7)	47
1. The term “interconnect means” designates structure.....	48
2. A “network” is a multi-hop interconnection between modules separated by one or more network nodes”	51
3. If “interconnect means” is subject to § 112, ¶ 6, the corresponding structure is a network	59
B. “mapping means” (claims 1 and 7)	61
C. The Preamble of Claim 6.....	62
D. District Court Constructions.....	63
1. “connection properties” (claims 1, 3-7)	63

2. “Integrated circuit” (claims 1-6)	64
3. “connection” (claims 1, and 3-7)	64
IX. ANALYSIS OF THE ’052 PATENT	65
A. Ground 1: Rădulescu2004 (Claims 1-2 and 5-7)	65
1. Rădulescu2004 Does Not Anticipate Independent Claim 6	65
2. Rădulescu2004 Does Not Anticipate Claims 1-2, 5, and 7.....	80
B. Ground 2: Purported Obviousness over Rădulescu2004 in view of OCP (Claims 1-7).....	80
1. A POSITA Would not Combine Rădulescu2004 and OCP	80
2. Rădulescu2004 in View of OCP Does Not Render Obvious Independent Claim 6	92
3. Rădulescu2004 in View of OCP Does Not Render Obvious Claims 1-5 and 7102	
C. Ground 3: Purported Anticipation by Weber (Claims 1-7).....	102
1. Weber Does Not Disclose a Single Integrated Circuit (Claims 1-7)	102
2. Weber Does Not Disclose the Claimed “Mapping Means” (Claims 1-5 and 7).....	110
3. Weber Does Not Anticipate Independent Claim 6	112
4. Weber Does Not Anticipate Claims 2-5, and 7.....	137
D. Ground 4: Purported Obviousness over Weber in view of Rădulescu2004 (Claims 1-7).....	137
1. Lack of Motivation to Combine Weber and Rădulescu2004	138
2. First Proposed Combination: Weber’s System of Communication Service Mapping with Rădulescu2004’s Network	148
3. Second Proposed Combination: Weber’s SoC with Rădulescu2004’s NoC and Network Interfaces.....	152
4. Motivation to Combine Weber and Rădulescu2004	138
X. CONCLUSION	158

I. INTRODUCTION

1. I have been retained on behalf of Patent Owner Network System Technologies, LLC (“NST” or “Patent Owner”), as an expert consultant.

2. As part of this work, I have been requested by counsel for NST to provide this declaration concerning the technical subject matter relevant to the *inter partes* review (“IPR”) petition of U.S. Patent No. 7,594,052 (“the ’052 patent”) filed by Qualcomm Incorporated (“Qualcomm”) and Samsung Electronics Co., Ltd. (“Samsung”). I am informed that Samsung has been terminated from this proceeding and therefore refer to Qualcomm as “Petitioner” herein.

3. I am being compensated for my time in the amount of \$700 per hour. I have no financial interest in NST or the ’052 patent, and my compensation is not contingent upon the substance of my declaration, any statements or opinions made, or the outcome of this proceeding.

4. I am informed and understand that the ’052 Patent was, or currently is, being asserted in district court litigation. I understand that the ’052 Patent is asserted against Petitioner in the Western District of Texas (“the Qualcomm Litigation”). I understand that the ’052 Patent was asserted in the Eastern District of Texas against Samsung (“the Samsung Litigation”). I also understand that the ’052 Patent was asserted against Texas Instruments in the Eastern District of Texas (“the TI Litigation”).

II. EXPERT QUALIFICATIONS AND BACKGROUND

5. I am currently a professor in the Department of Electrical Engineering and Computer Science at the University of California, Irvine. I have held this position since 2003.

6. Previously, I held the position of Professor and Chair of Electrical Engineering and Computer Science (ECE) at the University of California, Irvine, from 1998 to 2002. Prior to that, I was an Associate Professor in the ECE Department at the University of California, Irvine from 1993 to 1997. Prior to that, I was an Assistant Professor in the ECE Department at the University of California, Irvine from 1987 to 1992.

7. I earned a Bachelor of Science degree, with honors, in Electrical Engineering from the University of Texas at Austin in 1977. I earned a Master of Science degree in Electrical Engineering from the University of Texas at Austin in 1979, with a focus on communications theory and signal detection algorithms. I earned a Ph.D. in Computer Engineering from the University of Texas at Austin in 1987 with a focus on parallel processing and computer architecture.

8. I have over 45 years of experience in the Electrical Engineering and Computer Science field. That experience includes holding the position of Member of Technical Staff, No. 5 ESS Hardware/Software T1 Line Interface Design at AT&T Bell Labs from 1980 to 1984. In this position I received the Outstanding

Achievement award from AT&T Bell Labs in 1984. While at Bell Labs I developed T1 line framing and error monitoring algorithms that were stored on electrically erasable programmable read-only memory (EEPROM)—a charge based non-volatile memory which is considered the predecessor to contemporary flash memory. Additionally, I designed the microprocessor-based board that included the EEPROM among many other components for directly interfacing the T1 line.

9. I have particular experience in the development of System-on Chip (SoC) architecture with hundreds of cores (processors) based on the Network-on-Chip (NoC) technique. In multi-core processors, a NoC is responsible for data routing between different cores using basic networking principles. I have worked on various aspects of NoC design such as optimizing the efficiency of multi-core processing architectures by decreasing computing overhead and latency incurred from data transfer in a NoC using various routing policies and function configurations when running applications. The following paragraphs describe both my foundational and particular experience with this technology.

10. In 2002, I received the Best Paper award in Institute of Electrical and Electronics Engineers (IEEE) Transactions on Very Large Scale Integration (VLSI) Design related to task scheduling for multicore processors. I also received another Best Student Paper award co-authored with my doctoral students in the proceedings of Asia and South Pacific Design Automation Conference (ASPDAC)

'02. In 2020, I received the Best Paper award in the proceedings of the IEEE International Conference on Edge Computing (EDGE).

11. In 2014, I received the Khwarizmi International Award and became an IEEE Fellow. In 2015, I received Professor of the Year for the Department of Electrical Engineering and Computer Science at the University of California at Irvine.

12. Over the past 45+ years my research has focused on hardware, software, and VLSI design of low-power and high-performance multicore computer systems for wireless communication applications in mobile systems, computer graphics processing, 3D integrated circuits, memory structures (e.g., static random-access memory (SRAM), dynamic random-access memory (DRAM), flash memory, and register files), Network-on-Chip (NoC), and AI accelerators.

13. Since joining the University of California, Irvine in 1987, I have been teaching, researching, and consulting regarding almost all aspects of high performance computer systems, including but not limited to VLSI design, computer architecture, parallel implementation of computationally critical applications such as wireless communication signal processing and error correcting schemes, DRAMs, SRAMs, low power design, and interfacing multicores.

14. Over the course of my career, I have supervised more than 45 doctoral students and 30 masters' students who have taken key positions in the electrical

engineering and computer engineering sector. I have served as post-doctoral advisor for more than 10 scholars that worked with my research group.

15. In 2000, I cofounded a company called Morpho Technologies, which was focused on the design and development of low power and high-performance digital signal processors (DSPs) for mobile applications. As part of my duties as the chief technologist for the company, I evaluated patents, technical reports, and presentations related to memory chip designs, DSPs, and parallel processing algorithms for mobile platforms.

16. For Morpho, I designed, developed, and tested the key components of a Coarse Grain Reconfigurable Architecture (CGRA), including the memory block responsible for exchanging data between the host processor and the multicore processor (SIMD array processor). Morpho's main objective was to use a fully programmable platform based on the idea of CGRA to meet the power and performance requirements of a Software Defined Radio (SDR) for mobile embedded applications. In particular, I was able to show the efficiency of mapping an OFDM algorithm on the Morpho architecture.

17. I also designed the main data storage of the microcontroller (a reduced instruction set computer, or "Tiny RISC") that was used for storing temporary data between the Tiny RISC and the array processor, based on SRAM technology, which is a type of volatile memory. I was one of the primary contributors for the

prototype board design of the Morpho chip, which utilized DRAMs, field-programmable gate arrays (FPGAs), flash memory and other prototype board components.

18. As part of my teaching and research experience, I have been responsible for developing circuits, electronic modules, and related architectures for a variety of multicore systems. For example, I designed the memory block responsible for managing out-of-order execution of data registers corresponding to different threads in the multithreaded superscalar DSP that was developed at University of California, Irvine. I also designed the original temporary data storage subcomponents for the VLIW integer processor (ViPER) and the original Tiny RISC processor. ViPER was considered the very first ever VLIW processor designed using Metal Oxide Semiconductor Implementation Service (MOSIS) 1.2-micron technology that included memory and logic circuits. I have also done extensive work on the design of efficient data correction schemes using reconfigurable computing based on the Morpho parallel array architecture.

19. In 2006, I analyzed the hardware implementation of low-density parity-check (LDPC) as a consultant for a local flash memory module design company. In this capacity, with the help of one of my doctoral students we modeled the hardware and evaluated the computation complexity of using an application-specific integrated circuit (ASIC) for implementing LDPC on a flash memory-based board.

20. Throughout my career, I have worked on many aspects of computer architecture optimizations, encompassing hardware, processor allocation and scheduling on 2D NoC array, and software techniques. Ideas based on NoC design, modeling, and evaluation have been the hallmark of my research work for more than 15 years. I also extended my original ideas toward three-dimensional (3D) NoC and demonstrated some of the advantages of 3D integrated circuit (IC) design as well as potential complications that could arise from the thermal footprint of these architectures. NoC uses packet switching communication across a network of routers to facilitate core-to-core and core-to-memory communications. My main contributions have been on the critical aspects of NoC and IC design, including router design, through-silicon via (TSV) cross talk, power efficiency, task mapping and scheduling onto NoC, aging of ICs, and other related topics.

21. For NoC, I used Constraint Programming (CP) for solving the Voltage-Frequency Island (VFI) problem. In general, VFI is a superior design alternative in terms of thermal constraints, power consumption, and performance. Given a Communication Task Graph (CTG) and subsequent task assignments for cores, cores are allocated to the best possible places on the chip to minimize communication and power consumption. Moreover, an elegant CP model was proposed to solve the VFI problem by mapping and grouping cores at the same

time while scheduling the computation tasks as a limited capacity resource allocation model using realistic benchmark datasets.

22. I proposed a new architecture for NoC called diagonally linked mesh NoC, or DMesh NoC. This novel architecture not only enhances network transmission performance while maintaining a feasible implementation cost, but also provides a power-efficient solution for interconnection network scenarios. The proposed architecture uses an adaptive quasi-minimal routing algorithm so that it can improve average latency and saturation traffic load owing to its flexibility and adaptiveness. Based on these features, a congestion-aware routing algorithm was proposed to balance traffic load so as to alleviate congestion caused by high throughput network activities. Simulation results showed that saturation load is improved dramatically for various traffic patterns. Implementation results also suggested that employing diagonal links is a more area-efficient method for improving network performance than using large buffers. It was also shown that a congestion-aware router requires negligible cost overhead but provides better throughput.

23. I proposed a Quality of Service (QoS)-aware wireless NoC architecture that not only enables quality-oriented network transmission with a feasible implementation overhead but also maintains a well-balanced workload inside hybrid networks to enhance system throughput. The proposed advanced

routers employ congestion management which first grants high priority packets from more congested areas and then directs them to paths without contention. This adaptive load-balanced routing effectively relaxes congestion situations and prevents networks from excessive contention. Simulation results showed that average latency of high priority is significantly improved. Flexible buffer management further improved transmission performance and resource utilization efficiency.

24. I have worked on an efficient method for reaching the destination corresponding to 3D NoC nodes where power and latency have been considerably reduced as compared to what has been reported in the past. I devised an adaptive routing algorithm for partially connected 3D NoCs where an analytical formulation based on queuing theory for the first time was developed and its results were checked by simulation. I have worked on an energy and performance aware application mapping algorithm for inhomogeneous 3D NoCs that showed significant improvements. 2.5D NoC was the subject of another project that I did, for which I introduced a Controlled Load Balancing Mechanism (CLBM) for an efficient load balancing approach with more than 25% average latency and power reduction improvements with little accumulated overhead.

25. Through-silicon via (TSV) is one of the key enabling technologies for facilitating connections among cores of different layers of 3D ICs used for

implementing NoC architectures with 100's of cores. Some of my first work on designing routing algorithms for 3D NoC appeared in top IEEE journals where I introduced a new and robust method that is adaptive and can tolerate faults on the vertical TSV links. As part of this work, I proposed a novel algorithm to systematically generate hybrid 3D NoC topologies for a given application such that the vertical connections are minimized while the NoC performance is not sacrificed. Experimental results using synthetic and real-world traffic patterns show that the proposed algorithm generates optimized architectures with lower energy consumption and a significant reduction in packet delay as compared to the existing solutions.

26. Currently, I am working on developing accelerators and optimization techniques for implementation of efficient machine learning platforms. As part of this work, I am taking standard matrix vector computations that are the most demanding aspect of the Convolutional Neural Network (CNN) algorithms and implementing them on highly parallel structures, taking advantage of inherent parallelism of neural network layers. The focus of this work has been on developing low power inference engines for edge based AI nodes.

27. I have been involved in numerous research projects, journal articles, and books. The specific research projects, journal articles, and books may be found in my Curriculum Vitae, which is attached to my Declaration as Appendix A. Over

the years, I have also received funding accumulating over eight million dollars.

The specific funding and associated projects may be found in my Curriculum Vitae.

28. Additional details on my background and experience can be found in my *curriculum vitae*, attached hereto as Ex. 2071.

III. MATERIALS CONSIDERED

29. In forming the opinions set forth in this declaration, I have considered and relied upon my education, knowledge of the relevant field, and my experience. I have also reviewed and considered the '052 patent (Ex. 1001) and its file history (Ex. 1002), and at least the following additional materials:

- Petition for *Inter Partes* Review of the '052 Patent ("Petition")
- Declaration of Dr. Marwan Hassoun (Ex. 1003)
- Decision Granting *Inter Partes* Review of the '052 Patent (Paper 11) ("Institution Decision")
- A. Rădulescu, et al. "*An Efficient On-Chip Network Interface Offering Guaranteed Services, Shared-Memory Abstraction, and Flexible Network Configuration*," Proceeding of the Design, Automation and Test in Europe Conference and Exhibition, 2004 ("Rădulescu2004") (Ex. 1005)
- Open Core Protocol Specification, Release 1.0, OCP-IP Association (2001) ("OCP") (Ex. 1006)

- U.S. Patent Publication No. 2002/0129173 A1 (“Weber”) (Ex. 1007)
- “Networks on Chip; Axel Jantsch, Hannu Tenhunen (Eds.). Kluwer Academic Publishers, Boston (2003) (Ex. 1009)
- P. Guerrier et al., “*A generic architecture for on-chip packet switched interconnections*,” Proceedings Design, Automation and Test in Europe Conference and Exhibition 2000 (Ex. 1021)
- William J. Dally & Brian Towles, *Route Packets, Not Wires: On-Chip Interconnection Networks*, PROC. 38TH DESIGN AUTOMATION CONF. (2001) (Ex. 2005)
- Douglas Downing et al., Dictionary of Computer and Internet Terms, BARRON’S (7th ed. 2000) (Ex. 2015)
- Claim Construction Order, *Network System Technologies, LLC v. Texas Instruments Incorporated*, Civil Action No. 2:22-cv-00482-RWS (E.D. Tex.), Dkt. 150 (Jan. 3, 2024) (Ex. 2022)
- Defendants’ Joint Opening Claim Construction Brief, *Network System Technologies, LLC v. Qualcomm Incorporated et al.*, Civil Action No. 1:220-cv-1331-DAE) (W.D. Tex.), Dkt. 77 (Oct. 12, 2023) (Ex. 2040)

- K. Charles Janac, “Network-on-Chip (NoC): The Technology that Enabled Multi-processor Systems-on-Chip (MPSoCs),” *Multi-Processor System-on-Chip 1*, Wiley 2020 (Ex. 2072)
- Gul N. Khan, “NOC: Networks on Chip SoC Interconnection Structures,” Electrical, Computer & Biomedical Engineering Toronto Metropolitan University (Ex. 2073)
- U.S. Patent No. 7,373,449 (Ex. 2074)
- Deposition Transcript of Dr. Marwan Hassoun (Ex. 2068)

30. I have also considered the additional materials identified in the paragraphs below to the extent not specifically listed above.

IV. UNDERSTANDING OF LEGAL STANDARDS

31. I have a general understanding of validity based on my experience with patents, my career experience, and my discussions with counsel.

A. Prior Art

32. I understand that inventors may be entitled to a priority date earlier than an actual date of filing of a patent application that provides written description support for a particular claim to the extent that they can show a complete possession of such a particular claimed invention at such an earlier priority date, and reasonable diligence to reduce such a particular claimed invention to practice between such an earlier priority date and such an actual date of filing. I understand

that if Plaintiff contends that particular claims are entitled to an earlier priority date than the actual date of filing, then Plaintiff has the burden to prove this contention with specificity.

33. I understand that an invention by another must be made before the priority date of a particular patent claim in order to qualify as “prior art” under 35 U.S.C. § 102 or 35 U.S.C. § 103. I also understand that a printed publication or a product usage must be publicly available before the priority date of a particular patent claim in order to qualify as “prior art” under 35 U.S.C. § 102(a), that a printed publication or a product usage or offer for sale must be publicly available in the United States more than one year prior to the actual date of filing of a patent application that provides written description support for a particular claim in order quality as “prior art” under 35 U.S.C. § 102(b), and that the invention by another must be described in an application for patent filed in the United States before the priority date of a particular patent claim in order to qualify as “prior art” under 35 U.S.C. § 102(e). I understand that Petitioners have the burden of proving that any particular reference or product usage or offer for sale is prior art.

B. Anticipation

34. I have a general understanding of anticipation based on my experience with patents, my career experience, and my discussions with counsel.

35. I understand that the anticipation analysis is a two-step process. The first step is to determine the meaning and scope of the challenged claims. Each claim must be viewed as a whole, and it is improper to ignore any element of the claim. For a claim to be anticipated under U.S. patent law: (1) each and every claim must be identically disclosed, either explicitly or inherently, in a single prior art reference; (2) the claim elements disclosed in the single prior art reference must be arranged in the same way as in the claims; and (3) the identical invention must be disclosed in the single prior art reference, in as complete detail as forth in the claim. Where even one element is not disclosed in a reference, the anticipation contention fails. Moreover, to serve as an anticipatory reference, the reference itself must be enabled. In other words, the reference must provide enough information so that a person of ordinary skill in the art (“POSITA”) can practice the subject matter of the reference without undue experimentation.

36. I further understand that where a prior art reference fails to explicitly disclose a claim element, the prior art reference inherently discloses the claim element only if the prior art reference must necessarily include the undisclosed claim element. Inherency may not be established by probabilities or possibilities. The fact that an element may result from a given set of circumstances is not sufficient to prove inherency. I have applied these principles in forming my opinions in this matter.

37. I further understand that a dependent claim incorporates each and every limitation of the claim from which it depends. Thus, my understanding is that if a prior art reference fails to anticipate an independent claim, then that prior art reference also necessarily fails to anticipate all dependent claims that depend from that independent claim.

C. Obviousness

38. I have a general understanding of obviousness based on my experience with patents, my career experience, and my discussions with counsel.

39. I understand that a patent claim is invalid under 35 U.S.C. § 103 as being obvious only if the differences between the claimed invention and the prior art such that the subject matter as a whole would have been obvious at the time of invention to a POSITA. An obviousness analysis requires consideration of four factors: (1) scope and content of the prior art relied upon to challenge patentability; (2) differences between the prior art and the claimed invention; (3) level of ordinary skill in the art at the time of the invention; and (4) the objective evidence of non-obviousness, such as commercial success, unexpected results, the failure of others to achieve the results of the invention, a long-felt need which the invention fills, copying of the invention by competitors, praise for the invention, skepticism for the invention, or independent development.

40. I understand that a prior art reference is proper to use in an obviousness determination if the prior art reference is analogous art to the claimed invention. I understand that a prior art reference is analogous art if at least one of the following two circumstances is met. First, a prior art reference is analogous art if it is from the same field of endeavor as the claimed invention, even if the prior art reference addresses a different problem and/or arrives at a different solution. Second, a prior art reference is analogous art if the prior art reference is reasonably pertinent to problem faced by the inventor, even if it is not in the same field of endeavor as the claimed invention.

41. I understand that it must be shown that a POSITA at the time of the invention would have had a reasonable expectation that a modification or combination of one or more prior art references would have succeeded. Furthermore, I understand that a claim may be obvious in view of a single prior art reference, without the need to combine references, if the elements of the claim that are not found in the reference can be supplied by the knowledge or common sense of one of ordinary skill in the relevant art. However, I understand that it is inappropriate to resolve obviousness issues by a retrospective analysis or hindsight reconstruction of the prior art and that the use of such hindsight reconstruction is improper in analyzing the obviousness of a patent claim.

42. I further understand that the law recognizes several specific guidelines that inform the obviousness analysis. First, I understand that a hindsight reconstruction approach is not permitted. In other words, it is improper to use post-invention information to help perform the selection and combination of prior art references, and it is also improper to use the listing of elements in a claim as a blueprint to identify selected portions of different prior art references in an attempt to show that the claim is obvious.

43. Second, I understand that any prior art that specifically teaches away from the claim subject matter points to non-obviousness. In other words, prior art that would lead a POSITA to a specifically different solution than the claimed invention points to non-obviousness. Conversely, I understand that any prior art that contains any teaching, suggestion, or motivation to modify or combine such prior art references points to obviousness of such a modification or combination.

44. Third, while many combinations of the prior art may be “obvious to try,” I understand that any obvious-to-try analysis will not render a patent invalid unless it is shown that the possible combinations are: (1) sufficiently small in number so as to be reasonable to include that the combination would have been selected; and (2) such that the combination would have been believed to be one that would produce predictable and well understood results.

45. Fourth, I understand that if a claimed invention that arises from the modification or combination of one or more prior art references uses known methods or techniques that yield predictable results, then that factor also points to obviousness.

46. Fifth, I understand that if a claimed invention that arises from the modification or combination of one or more prior art references is the result of known work in one field prompting variations of it for use in the same field or a different one based on design incentives or other market forces that yield predictable variations, then that factor also points to obviousness.

47. Sixth, I understand that if a claimed invention that arises from the modification or combination of one or more prior art references is the result of routine optimization, then that factor also points to obviousness.

48. Seventh, I understand that if a claimed invention that arises from the modification or combination of one or more prior art references is the result of a substitution of one known prior art element for another known prior art element to yield predictable results, then that factor also points to obviousness.

49. I further understand that a dependent claim incorporates each and every limitation of the claim from which it depends. Thus, my understanding is that if a prior art reference fails or a combination of prior art references fail to render obvious an independent claim, then that prior art reference fails or that combination

of prior art references fail to render obvious all dependent claims that depend from the independent claim.

D. Claim Construction

50. I understand that claim construction is the process by which a court determines the scope and meaning of terms used in the claims of a patent. I understand that the goal of this process is to give claim terms the ordinary and customary meaning they would have had to a person of ordinary skill in the art (“POSITA”) at the time of the invention, after reading the entire patent and its prosecution history.

51. I understand that it is possible that the patent specification may reveal a special definition given to a claim term by the patentee that differs from the meaning it would otherwise have to a POSITA. In such cases, I understand that the patentee’s definition usually controls.

52. I understand that the prosecution history of a patent can inform the meaning of some claim language and must be taken into account in construing the claims.

53. I understand that, in some cases, the court may consider extrinsic evidence, such as dictionaries, treatises, and expert opinions, to understand the underlying technology and the way in which claim terms would be understood by a POSITA at the relevant time.

54. I understand that as a general rule, unless the steps of a method actually recite an order, the steps are not ordinarily construed to require one. However, I understand that such a result can ensue when the method steps implicitly require that they be performed in the order written. I understand that this determination involves a two-part inquiry: (1) whether the claim language requires an order as a matter of logic or grammar; and (2) if the answer to (1) is in the negative, whether the specification directly or implicitly requires such a narrow construction.

55. I understand that 35 U.S.C. § 112, ¶ 6 creates an exception to the general rules of claim construction called a “means-plus-function” limitation. I understand that a means-plus-function limitation describes what the limitation does (i.e., its function), opposed to what it is (i.e., its structure). I also understand that a means-plus-function limitation is construed to cover only the corresponding structure, material, or acts described in the specification for performing the function, and equivalents thereof.

56. I understand that a limitation is presumed to be a means-plus-function limitation if (i) the claim limitation uses the phrase “means for,” (ii) the phrase “means for” is modified by functional language, and (iii) the phrase “means for” is not modified by sufficient structure for achieving the specified function. I understand that to determine whether a claim recites sufficient structure, it is

sufficient if the claim term is used in common parlance or by persons of skill in the pertinent art to designate structure, even if the term covers a broad class of structures and even if the term identifies the structures by their function.

57. I understand that in determining whether a claim term invokes § 112, ¶ 6, the essential inquiry is not merely the presence or absence of the word ‘means’ but whether the words of the claim are understood by persons of ordinary skill in the art to have a sufficiently definite meaning as the name for structure.

58. I understand that interpreting a means-plus-function limitation is a two-step process. First, the claimed function must be identified. Second, the corresponding structure in the specification, if any, that corresponds to the claimed function(s) must be identified. I understand that structure disclosed in the specification is corresponding structure only if the specification or prosecution history clearly links or associates the structure to the function recited in the claims. I also understand that where there are multiple claimed functions, the specification must disclose adequate corresponding structure to perform all of those functions. I further understand that the claim is indefinite if the specification fails to disclose adequate corresponding structure.

V. PERSON OF ORDINARY SKILL IN THE ART

59. I understand that, in assessing the level of ordinary skill, one may consider several factors, including: (1) the education level of the inventor; (2) the

education level of active workers in the field; (3) the types of problems encountered in the art; (4) prior art solutions to these problems; (5) the rate of innovation in the field; and (6) the sophistication of the technology.

60. I understand that the District Court in the TI Litigation adopted the following definition for a person of ordinary skill in the art (“POSITA”): “a POSITA would have at least a Bachelor’s degree in electrical engineering, computer engineering, or a related field, and at least two years of work experience in the field of SoC and NoC technologies.” Ex. 2022.

61. I understand that Petitioner and Dr. Hassoun propose the following level of ordinary skill in the art for this IPR:

The relevant POSA for the Accused Patents would typically have had at least either a Master’s degree in Electrical Engineering, or a related field, as well as at least three years of work experience in the field of SoC design, or a Bachelor’s degree in Electrical Engineering, or related field, as well as at least five years of work experience in the field of SoC design. This POSA would have had an understanding of SoC interconnects, as well as the concepts of large-scale networks.

Ex. 1003 at ¶ 32.

62. I understand that in the Institution Decision, the Board agreed with and adopted Petitioner and Dr. Hassoun’s proposed level of skill in the art for the ‘052 Patent. Paper 11 at 10.

63. I met and exceeded the level of skill articulated in the TI Litigation and the level of skill articulated and adopted by the Board in the Institution Decision. I am also very familiar with individuals having that articulated level of skill. I have applied Petitioner's definition of a POSITA throughout this declaration. To the extent the Board were to determine that the level of skill articulated in the TI Litigation is appropriate, that would not change any of my opinions.

VI. TECHNOLOGY OVERVIEW

A. Technical Background

64. A system-on-chip ("SoC") is an integrated circuit in which the components of a system are integrated on the chip, including memories, processors, interfaces, and logic. Each SoC must include an interconnect, which interconnects the components of the system so that they may communicate with each other.

65. In the early years of computers, the components of a system were directly interconnected using point-to-point links. For example, the first computer system was named ENIAC (Electrical Numerical Integrator and Computer) and was developed in the United States during World War II. ENIAC used point-to-point links to connect the components of the systems such as accumulators, central control, and I/O devices.

66. The industry soon thereafter realized a foundational problem with point-to-point link interconnects: all the components in the system must be directly connected with such links. This resulted in increased hardware overhead, especially as the number of components included in the system increased over time.

67. An advancement over the traditional point-to-point link interconnects is what's been termed a "bus" interconnect. The formal concept of a bus, where a single communication pathway could be shared by multiple devices, evolved in the 1950s and 1960s. For example, the UNIVAC 1 (1951) and IBM 7090 (1959) made great progress towards the goal for managing data flow between the central processor, memory, and peripherals. The development of more standardized bus architectures, such as the S-100 bus (Altair bus) was developed in the 1974. As the S-100 bus gained in popularity, it became necessary to introduce a formal bus specification to ensure compatibility and interoperability among products from different manufacturers. There was also a need to extend the bus capability so that it could support more complex processors. In 1979, Standard Specification for S-100 Bus Interface Devices was introduced. This need for standardization led to the development of the IEEE 696 specification, which was formally adopted by the IEEE in 1983. The IEEE 696 bus specification was a critical early standard in the personal computing world because it provided a foundation for hardware compatibility and expansion during the early years of the microcomputer industry.

68. In modern times, as the number of computational, functional, and memory units on a single chip increased, the need for communication between those units also increases. It is therefore important to choose an appropriate interconnect scheme that is scalable, efficient, reliable, and flexible. A scalable interconnect is an interconnect in which additional or fewer computational, functional, and/or memory units could be included in the SoC without issue. An efficient interconnect is an interconnect that consumes relatively low silicon area and power while offering high bandwidth and low latency. A flexible interconnect is an interconnect in which it is easy to attach or remove a component or an IP from the interconnect without modifying the interconnect completely.

69. SoC designers have the option of using a point-to-point interconnect, single-bus architecture, hierarchical bus architecture, or a network-on-chip interconnect. Point-to-point interconnects are the least scalable with limited flexibility and are optimal when future design changes are not needed and are usually reserved for high performance applications. For example, a point-to-point connection may exist between two components that are tightly coupled and need a high bandwidth and low latency interconnect to meet power and performance requirements.

70. A bus interconnect is a shared communication interconnect and is used to connect various components or modules within a device. This design

approach allows different IPs, cores, memory, and peripheral devices to communicate with each other by transmitting data, addresses, and control signals over the shared bus. Bus-based interconnects have been the dominant form of interconnect prior to the invention described in the '052 Patent and continue to be widely used when the number of integrated components is in the order of tens. A bus is essentially a set of parallel wires on the chip that carry data, address, and control signals between different components. In a bus, one component (i.e., the master) initiates communication while the other component (i.e., the slave) responds. In bus-based interconnects, only one master can have control of the bus at a time. Thus, when there are multiple master components connected to the bus, arbitration logic must be employed to dictate which master has control of the bus at any given time to avoid conflicts. Each component connected to the bus typically has a bus interface unit that manages the interaction between the device and the bus, ensuring that data is transmitted correctly and provides seamless interaction with the bus communication protocol.

71. While buses solved the some of the limitations issues associated with point-to-point interconnects, a bus is a bottleneck when the number of components connected to the bus increases. This is because all the components share the same bus and there can be contention for bus access, leading to bottlenecks and significant delays in bus response time. The time it takes for data to travel along

the bus can introduce delays, particularly in long buses or when multiple devices are trying to communicate simultaneously. The need to drive signals across a shared bus can lead to increased power consumption, particularly in high performance systems.

72. Thus, as the complexity and number of components on a chip have increased, traditional bus-based designs have faced scalability challenges. This led to the development of NoC architectures, where a network of interconnected switches and routers (and potentially traditional interconnects) replaces the traditional bus, providing better performance and scalability for modern SoCs.

73. NoCs represent an advanced communication architecture for connecting the ever-increasing number of components on a chip. Unlike traditional bus-based designs which rely on a single shared communication highway, NoCs implement a network of interconnected routers and links to facilitate data transfer between the various components on a chip. Routers and switches are the main components in a NoC architecture because they handle the data packets coming from neighboring routers and switches, determine the best path, and forward the packet to their destination. Links are the physical and logical connections between routers and switches and are implemented as wires or interconnects within the chip, thereby facilitating packet transmission for one router or switch to another.

74. Unlike bus-based systems, NoCs can efficiently handle a large number of cores and IP blocks, making NoCs suitable for modern SoCs. NoCs allow multiple data transfers to occur simultaneously across different parts of the network, improving overall system performance and power consumption. However, designing a NoC is more complex than a traditional bus-based interconnect because NoCs require sophisticated routing algorithms, flow control, load balancing, and error handling.

75. NoCs are sophisticated communication architectures that provide an efficient, scalable, and high-performing means of transferring data within a SoC. By using a network of routers, switches, and link, NoCs overcome the limitations of traditional bus-based interconnects, support the growing demand for parallelism, scalability, and energy efficiency in modern computing systems. NoCs are a critical technology in the development of modern electronic devices.

B. Overview of the '052 Patent

76. I understand that Petitioners challenge claims 1-7 of the '052 Patent in this IPR.

77. Independent claim 1 is reproduced below:

1. [1pre] Integrated circuit, comprising:

[1a] 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,

[1b] an interconnect means (N) for coupling said plurality of processing modules (M, S) and for enabling a connection based communication having a set of connection properties,

[1c] at least one network interface (NI) associated to said at least one first of said processing modules for controlling the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N), and

[1d] a mapping means (A) for 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 said 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).

78. Independent claim 6 is reproduced below:

6. [6pre] 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:

[6a] coupling said plurality of processing modules (M, S) by an interconnect means (N) and

[6b] enabling a connection based communication having a set of connection properties,

[6c] controlling the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N) by at least one network interface (NI) associated to said at least one first of said processing modules,

[6d] 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 said at least one communication service identification.

79. As the '052 Patent explains, systems-on-chip ("SoCs") "show[ed] a continuous increase in complexity due to the ever increasing need for implementing new features and improvements of existing functions" at the time of the invention. Ex. 1001 at 1:14-16. Conventionally, the modules on the SoC "communicate to each other via a bus." *Id.* at 1:25-27. However, buses were proving "no longer practical" because "the large number of modules form[ed] a too high bus load" and "the bus form[ed] a communication bottleneck as it enable[d] only one device to send data to the bus." *Id.* at 1:27-32. The '052 Patent describes a network-on-chip ("NoC") interconnect which provides "an effective way to overcome these disadvantages" of bus-based SoCs. *Id.* at 1:33-34. NoCs are beneficial for two separate and distinct reasons: (1) NoCs "share wires, lowering their number and increasing their utilization" and "are scalable compared to buses" (*id.* at 1:37-42) and (2) NoCs "decouple computation from communication" by "using protocol stacks, which provide well-defined interfaces separating communication service usage from service implementation" (*id.* at 1:42-47).

80. Decoupling computation from communication is a significant advancement of NoC interconnects compared to previous interconnects. “This is because, in contrast to existing on-chip interconnect (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.” Ex. 1001 at 1:50-54; *see also id.* at 2:31-42 (“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. This is in contrast to existing interconnects (i.e., buses) where modules are directly connected.”).


81. The '052 Patent also recognizes the importance of “reduc[ing] the cost of system design through re-use of applications and architectures.” Ex. 1001 at 2:43-45. For example, “the IP blocks or modules” which communicate “in modern on-chip communication protocols (e.g., Device Translation Level DTL, Open Core Protocol OCP, and AXI-Protocol), [] must be incorporated to the network on chip.” *Id.* at 2:65-3:2. To re-use architectures (i.e., the modules and functional blocks on existing bus-based SoCs), the '052 Patent teaches that the interconnect “must implement diverse application behaviors with application-dependent IPs in an application-independent manner.” *Id.* at 2:55-57. “In other words, the NoC services

largely define the platform, and it naturally provides differentiated service by means of a ... protocol stack.” *Id.* at 2:59-62.

82. Therefore, “[t]he 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.” Ex. 1001 at 4:15-18.

83. The ‘052 Patent accomplishes its goal of offering differentiated services for protocols by converting the IP modules’ communication protocols to the communication protocol of the network. For example, the ‘052 Patent describes that a master module M requests a communication service to a slave module S. Ex. 1001 at 5:53-54. This request is sent in a first protocol (e.g., OCP). The slave S receives the request in its protocol and performs the requested action. *Id.* at 5:54-58. However, because the master module M and the slave S are not directly connected, the master module M’s request must first go through the network. The ‘052 Patent provides that “adapter units A serve to connect the IP modules and their communication protocol to the network interface which in turn serve to connect to the network N.” *Id.* at 5:58-61. “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.” *Id.* at 5:61-66. Once the message is received by the network interface associated with the slave module S, the adapter unit A within

that network interface would similarly perform translation; however, because the network interface receives the message from the network as opposed to the master module M, the network interface would translate packet into the message msg so that it may be sent to the adapter unit A, which then translates the message msg into the communication protocol of the slave module S. This protocol conversion concept is shown with reference to FIG. 1, reproduced below.




EX1001, FIG. 1

84. As shown in FIG. 1, the master module M issues a request for a communication service to a slave module S. The requested communication service is issued in the DTL protocol. The requested communication service is sent to the adapter unit A, which translate the communication service in the DTL protocol to a message msg in a separate protocol. Ex. 1001 at 5:58-66. The message msg is then sent to the network interface NI which translates the message msg into flits. *Id.* The flits are sent over the multi-hop network N until they arrive at the network interface associated with the slave module S. The network interface NI associated

with the slave module S converts the flits to a message msg. The message msg is sent to the adapter unit, which converts the message msg into slave module S's communication protocol, which, in this example, is the MTL protocol. The slave module S is now able to receive the message in its native protocol and perform the action requested by the master module M.


85. The adapter unit A, which translates the requested communication service into messages msg, is also called “mapping means” or shell in the specification. Ex. 1001 at 3:22-27, 6:41-46. The adapter unit A “comprise[s] a sequentializing unit seq and a desequentializing unit deseq.” *Id.* at 6:46-48. The structure of the adapter unit is shown in Figures 2A and 2B of the '052 Patent.



Ex. 1001 at FIG. 2A (annotated)

86. The '052 Patent discloses that the “adapter unit A may also be arranged in the network interface.” Ex. 1001 at 5:10-13. Further, 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.” Ex. 1001 at 6:30-33. “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.” *Id.* at 6:33-36. An example of the network interface with adapter units is shown in FIG. 3, reproduced below.




Ex. 1001 at FIG. 3 (annotated)

87. In FIG. 3 above, from right to left, the “network interface NI provides a conversion of the packet-based communication of the network N to the higher-level protocol,” in this example AXI and DTL, “that the IP modules use.” Ex. 1001 at 6:61-64. The network interface allows the NoC interconnect to “offer[] high-

level services, such as transaction ordering, throughput and latency guarantees, and end-to-end flow control based on connections through the network” because the “NI kernel provides” the aforementioned high-level services and the adapter units “implement additional functionality, such as multicast and narrowcast connections, and adaptors [sic] to existing protocols, such as AXI or DTL.” *Id.* at 7:25-39.

88. The protocol conversion is based on mapping requested communication service to a connection based on the identification means of existing protocols (e.g., AXI, DTL, or OCP). The identification means may be a communication thread or an address range. Ex. 1001 at 5:66-6:3. In an example disclosed in the ‘052 Patent, “network interface NI associated to the master M may map the communication to the first slave S1, i.e., to the address range 0000-7FFFF, to a connection with the connection id cid=0 and the communication to the second slave [sic] S2, i.e., to the address range 8000-FFFF, to a connection with the connection id cid=1.” *Id.* at 8:15-20. Even if a master’s communication protocol does not support communication threads, “address ranges can be mapped to specific connections.” *Id.* at 8:21-23. “An address range can identify one or more slave modules, or a memory region within a slave module.” *Id.* at 8:23-25. An exemplary mapping according to an address range is depicted in Figure 4, reproduced below.



Ex. 1001 at FIG. 4

89. The ‘052 Patent also discloses that the mapping a request to a connection may be done according to communication threads. Ex. 1001 at 7:40-44. “Where communication threads are used to identify traffic classes of a communication based on specific communication parameters, such a thread and the associated traffic class can be mapped to a connection based on specific corresponding communication properties.” *Id.* at 7:44-48. In other words, the mapping recited in the ‘052 Patent is dependent on the corresponding connection properties and knowledge of the system. For example, the network and network interface knows whether the request is destined for a fast memory or a slow memory and would map the request accordingly, i.e., “[t]he thread identifying the fast traffic class can be mapped to a connection configured to guarantee high

bandwidth, while the slow traffic class can be mapped to a connection with a low guaranteed bandwidth.” *Id.* at 7:51-55.

90. Whether the system is mapping requests to connections according to an address range or according to a thread, the system always considers the “connection data” associated with the particular connection. Ex. 1001 at 8:63-65. “The connection data comprises specification for data transport such as best effort or guaranteed through, the maximum allowed jitter, the maximum allowed latency, flow control Y/N, or ordering Y/N.” *Id.* at 8:65-9:2. “If a message is sent to a particular address range in the system, or belongs to a particular thread, a connection is assigned to assure that the message is transmitted via a particular set of transmission requirements.” *Id.* at 9:2-5. Therefore, the ‘052 Patent teaches that the disclosed mapping ensures selection of the right connection for a particular request.

91. This type of protocol conversion and mapping is critical “for communication protocols which do not include the notion of connection, but are implemented with an interconnect supporting connections.” Ex. 1001 at 8:47-50. In other words, the protocol conversion and mapping allows for modules communicating in a transaction-based communication protocol such as AXI, DTL, or OCP (i.e., protocols which do not include the notion of connection) to communicate over a network that transmits data in a network-based

communication protocol. *See id.* at 8:54-58 (“Backward compatibility of on-chip networks is achieved, by means of the adaptor units/shells to enable existi[ng] protocols, such as DTL, AXI and OCP, to use differential services as offered by a connected based interconnect, such as an on-chip network.”).

VII. OVERVIEW OF PETITIONER’S REFERENCES

A. Rădulescu2004 (Ex. 1005)

92. Rădulescu2004 describes a “network interface” that “offer[s] shared-memory abstraction, which is independent of the network implementation.” Ex. 1005 at Abstract. The network interface comprises two parts: “(a) the *NI kernel*, which implements the channels, packetizes messages and schedules them to the routers, implements the end-to-end flow control, and the clock domain crossing, and (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.” Ex. 1005 at 3 § 4. The network interface is shown in Figure 1, reproduced below.





Figure 1. NI kernel and shells

Ex. 1005 at FIG. 1

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.” Ex. 1005 at 3 § 4.1.

93. The shells implement narrowcast connections which “are connections between one master and several slaves, where each transaction is executed by a single slave based on the address provided in the transaction.” *Id.* at 4 § 4.2. The narrowcast shell is shown in Figure 3, reproduced below.


**Figure 3. Narrowcast shell**

Ex. 1005 at FIG. 3

“Within a narrowcast connection, the slave for which the transaction is destined is selected based on the address (CONN clock). The address range assigned to a slave is configurable in the narrowcast module.” Ex. 1005 at 4 § 4.2. Rădulescu2004 provides for in order response and request delivery. *Id.* (“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. In-order delivery per slave of request messages is already provided by the point-to-point connections.”).

B. OCP (Ex. 1006)

94. OCP (Open Core Protocol Specification) is a bus-based communication protocol. Ex. 1006 at 13. OCP describes a simple wrapped bus and three IP cores in Figure 1, reproduced below. Ex. 1006 at 14.



Ex. 1006 at FIG. 1

OCP's "system initiator (as the OCP master) presents command, control, and possibly data to its connection slave." Ex. 1006 at 13. "Each OCP interface on the core must ... [c]ontain at least the basic OCP signals." Ex. 1006 at 15. OCP's basic signals include Clk, MAddr, MCmd, MData, SCmdAccept, SData, and SResp. *Id.* at 24. The MAddr signal "specifies the slave-dependent address of the resource targeted by the current transfer." *Id.*


95. OCP also describes a series of "**optional** dataflow signals [that] are grouped into simple and complex extensions." Ex. 1006 at 24. OCP's optional simple signals are: MAddrSpace, MBurst, MByteEn, MDataValid, MRespAccept, and SDataAccept. *Id.* at 26. The optional simple MAddrSpace signal "is an extension of the MAddr field and is used to indicate the address region of a transfer." *Id.* OCP's optional complex extensions are: MConnID, MDataThreadID,

MThreadBusy, MThreadID, SThreadBusy, and SThreadID. *Id.* at 27. The optional complex MThreadID signal is the “[r]equest thread identifier” and “provides the thread identifier associated with the current transfer request.” *Id.* at 28. A basic OCP request only includes the required signals and could include the optional signals “to support additional core communication requirements.” *Id.* at 23.

C. Weber (Ex. 1007)

96. Weber describes a communication system designed to “meet the requirements of today’s large, complex electronics systems” which failed to “support true concurrency of operations” and “provide[d] limited capability to preserve order in one transaction stream.” Ex. 1007 at ¶¶ [0007]-[0009]. Weber’s communication system was designed to meet today’s requirements while also “decreas[ing] physical size and cost per function” and maintain a low “form-factor” of the communication system. *Id.* at ¶¶ [0002]-[0003].

97. Weber describes a shared-bus system that allows subsystems in “large, complex electronics systems to interoperate efficiently regardless of their varying performance characteristics and requirements.” Ex. 1007 at ¶ [0009]. Weber’s Figure 1, reproduced below, illustrates several functional blocks on separate integrated circuits or chips (102, 104, 106, 108, 110) that are interconnected via shared buses 112 and 114. Ex. *Id.* at ¶ [0037].



Ex. 1007 at FIG. 1

Weber's interfaces (e.g., 120) implement communication within the system. Ex. 1007 at ¶ [0045]. The interfaces are directly to each functional block and "decode[] an address on shared communication bus 814 to determine if a write is to be performed." *Id.*

98. In a separate embodiment, Weber's Figure 14, reproduced below, describes a shared-bus approach. Ex. 1007 at ¶ [0095].




FIG. 14

Ex. 1007 at FIG. 14

In Figure 14, functional blocks (e.g., 1405) are interconnected by communication subsystems (e.g., 1440). *Id.* at ¶ [0095]. A “data flow” between the functional blocks and the communication subsystems exists where “[a] communication subsystem transmits data items from one of its input ports to an output port, using a dedicated data channel for each data flow.” *Id.* at ¶ [0096].

VIII. CLAIM CONSTRUCTION

99. I understand that the following terms of '052 Patent have been proposed for construction by either Petitioner or the Patent Owner, or have been construed, either by agreement or by the Court in related district court litigation. For all other claim terms, I have applied the plain and ordinary that a POSITA would understand for such terms at the time of the invention.

A. “interconnect means” (claims 1, 6, and 7)

100. I understand that Petitioner and Patent Owner have proposed the following constructions for the term “interconnect means,” which appears in all challenged claims of the ‘052 Patent.

Term	Petitioners’ Proposed Construction	Patent Owner’s Proposed Construction
“interconnect means”	Subject to § 112, ¶ 6 Function (claim 6): coupling said plurality of processing modules (M, S). Function (claims 1 and 7): the function of claim 6 plus enabling a connection based communication having a set of connection properties Structure: any of a bus, switch, point-to-point wire, or network. Petition at 10.	“multi-hop interconnection between modules separated by one or more network nodes” If subject to § 112, ¶ 6 Function (claims 1, 6, and 7): coupling said plurality of processing modules (M, S). Function (claims 1 and 7): enabling a connection based communication having a set of connection properties Structure: multi-hop interconnection between modules separated by one or more network nodes

101. I understand that in litigation involving Patent Owner and Texas Instruments in the Eastern District of Texas, the court construed the term “interconnect means” as “multi-hop interconnection between modules separated by one or more network nodes.”

102. In this proceeding, I understand that there are two disputes for the term “interconnect means.” The first dispute is whether the term “interconnect means” recites sufficient structure or whether it is subject to means-plus-function treatment under § 112, ¶ 6. If this term is subject to § 112, ¶ 6, the parties agree on the corresponding function. The second dispute is what is the appropriate corresponding structure for the “interconnect means” if this term is subject to § 112, ¶ 6. My opinions on these issues follow.


1. The term “interconnect means” designates structure

103. The term “interconnect means” as used in the ‘052 Patent would be understood by a POSITA to have a sufficiently definite meaning as the name for structure. This is because the specification uses the terms “interconnect means” and “network” interchangeably. A POSITA would understand that a “network” is a structure. I also understand the Board’s Institution Decision found that the corresponding structure for the “interconnect means” is a network. Institution Decision at 18.

104. Each of claims 1, 6, and 7 recite an “interconnect means (N),” “processing modules (M, S),” and a “network interface (NI).” Ex. 1001 at 9:25-10:58 (emphasis added). A POSITA would understand that the “(N)” suffix following the term “interconnect means” refers to the exemplary “network” described in the specification.

105. For example, FIG. 1 of the '052 Patent illustrates “two master modules M and a slave module S each with an associated network interface NI.”


Ex. 1001 at 5:6-8.



Ex. 1001 at FIG. 1 (annotated)

FIG. 1 also illustrates a “network N” that “may be comprised of a plurality of interconnected network routers R.” Ex. 1001 at 5:24-25.

106. Similarly, FIG. 4 of the '052 Patent depicts that “[t]he network on chip comprises a master M and two slaves S1, S2 each being connected to a network via a network interface NI, respectively.” Ex. 1001 at 8:9-12.



Ex. 1001 at FIG. 4

107. By virtue of including “(N)” after “interconnect means,” the claim language itself equates the claimed “interconnect means” to the “network” described in the specification.

108. Both FIG. 1 and FIG. 4 show that the network N interconnects modules. *See* Ex. 1001 at FIG. 1, FIG. 4. The same function is ascribed to the interconnect means. *Id.* at Abstract (“Furthermore, an interconnect means (N) is provided for coupling said plurality of processing modules (M, S) and for enabling a connection based communication having a set of connection properties”); *id.* at 1:7-8 (“The intervention relates to an integrated circuit having a plurality of processing modules and an interconnect means for coupling said plurality of processing modules.”); *id.* at 3:51-54 (“Said plurality of processing modules are

coupled by an interconnect means and enabling a connection based communication having a set of connection properties.”).

109. The specification also repeatedly refers to the “interconnect means” or “interconnect” as a network. For example, the specification teaches that “[t]he interconnect is embodied as a network on chip NOC, allowing communication between all aforementioned modules.” Ex. 1001 at 4:46-50. The Description of the Invention section describes networks as the interconnect means. Ex. 1001, 5:7-31, 6:56-64, 7:25-28, 8:9-12, 8:54-58.

110. Thus, a POSITA would understand that the “interconnect means” recited in the claims of the ‘052 Patent refers to a network. A POSITA would also understand that “network” is a structural term. Thus, the term “interconnect means” in the claims of the ‘052 Patent conveys structure to a POSITA.

2. A “network” is a multi-hop interconnection between modules separated by one or more network nodes”

111. In the context of the ‘052 Patent, a POSITA would understand that a “network” (and thus an “interconnect means”) is a multi-hop interconnection between modules separated by one or more network nodes.

112. The specification confirms that a “network” is a “multi-hop interconnect between modules separated by one or more network nodes.” Specifically, the ‘052 Patent provides that “[i]ntroducing networks as on-chip

interconnects radically changes the communication when compared to direct interconnects, such as buses or switches ... 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 at 2:31-36 (emphasis added). Thus, the specification distinguishes the claimed interconnect or network from conventional interconnects like buses or switches because of its multi-hop nature. In view of these teachings in the specification, a POSITA would not understand the “interconnect means” in the claims read onto the conventional interconnects that are expressly distinguished and disparaged.


113. Because a POSITA would understand that “interconnect means (N)” is used interchangeably with “network,” and it defines as a “network” as a “multi-hop interconnection separated by one or more network nodes,” a POSITA would understand that an “interconnect means (N)” is a “multi-hop interconnection separated by one or more network nodes.”

114. A “multi-hop network” (i.e., a “interconnect means”) is a type of interconnect where data is transmitted from a source to a destination through one or more intermediate routers or switches, or “hops.” A POSITA would understand that a “hop” is a router/switch that a packet originating from an arbitrary protocol may traverse. *See* Ex. 2005 at 4 (“At each hop, the packet moves from one link to another.”). Each router or switch in the network forwards data to the next router or

switch until the data reaches its final destination. A POSITA would understand that the routers or switches in the data path are not required to be directly connected.

115. A POSITA would understand that routers and switches are both networking components that serve similar functions. Both routers and switches connect functional blocks to the network and transmit data from said functional blocks across the network. The difference between routers and switches comes down to scale: a switch can connect functional blocks via a single connection to the network while a router can connect functional blocks and switches across multiple connections to the network. A POSITA would understand that a “hop” is a router or switch that a packet traverses.

116. As shown in the exemplary figure below, a router or switch is **not directly** connected to every other router or switch in the network; instead, each router or switch is indirectly connected to the remaining routers or switches through the “network” of routers and switches. Each router is directly connected to its immediate neighbors, and indirectly connected to the other routers in the network.




Ex. 2073 at 30 (annotated)

A POSITA would understand that the above figure is a representation of a network topology called a “2D mesh.” As can be seen from the figure above, router/switch 1 is directly connected to router/switch 2 and router/switch 5 but is indirectly connected to the remaining routers/switches. If data from router/switch 1 needed to be sent to router/switch 10, the data would traverse multiple “hops” in the network (i.e., from router/switch 1 to router/switch 5, from router/switch 5 to either router/switch 6 or router/switch 9, and finally from router/switch 6 or router/switch 9 to router/switch 10). The routers and switches of the network make a routing decision as to which router or switch to forward the data to. In the above example, router/switch 1 decided to send the data to router/switch 5 instead of router/switch 2. As evidenced by the optioned routing path above, the routers/switches

selectively transmit data based on numerous factors, including network resources, the connection properties of the connection, and efficiency.

117. A multi-hop network is very different than a bus, where sources and destinations are connected to the same structure (i.e., the bus).

Buses to NoC



Ex. 2073 at 15

As shown in the figure above, each functional unit is directly connected to a shared-bus and the routing path in the shared-bus is fixed. For example, in the exemplary bus structure, the initiating functional unit is directly connected to a shared bus. A request sent from the initiating functional unit would be received by the shared bus. A request sent from the initiating functional unit would be received by the shared bus at an input port, the data would traverse via a wire from the input port to a corresponding output port, and, upon exiting the bus, the data would be sent to the target functional unit, which is directly connected to the output port of the bus. Therefore, buses do not require the use of routers or switches, making

them simpler, but also insufficient to carry out large-scale transactions with a large quantity of sources and destinations. Thus, as the number of functional units connected to the bus increase, there is a greater chance of congestion within the bus, which negatively impacts data transfer throughout the integrated circuit.

118. Another fundamental aspect of a multi-hop network is that it transports “data packets” (or “packets”) along each hop. *See* Ex. 2005 at 4 (“At each hop, the packet moves from one link to another.”); *see also* Ex. 2073 (“packet ... keeps hopping from router to router ...”). A packet is a bundle of data that is transmitted over a network according to specific protocols, some of which include TCP/IP (specific to computer networks, i.e., the Internet) and NTTP (specific to NoCs, i.e., on-chip networks).

119. I note that Dr. Hassoun agreed that in “[a] network on a chip, packets are transmitted over the network.” Ex. 2068 at 26:22-24; *id.* at 27:3-5 (“the packets can travel from one network node to a second network node.”). Dr. Hassoun also agreed that a hop is a router or switch that a packet may traverse. *Id.* at 27:17-21 (“So if the network had a router and the packet traveled from a first network node to the router and then from the router to a second network node, is that a hop? A. Yes, the router would be a hop.”).

120. A bus is incapable of transmitting packets because the bus communicates in the same protocol as each functional unit. Examples of such


functional unit/bus protocols, as discussed in the 052 Patent, are AXI, OCP, DTL. *See e.g.*, Ex. 1001 at 6:30-38. What makes the on-chip multi-hop network (i.e., the “interconnect means”) unique is that it only transmits packets and does not transmit data existing in the communication protocol of the functional units. Thus, before the data is sent through the network, it must be converted from the communication protocol of the functional units (e.g., AXI, OCP, DTL) to a protocol suitable for network transmission (e.g., NTTP). In other words, the data must be “packetized.” For example, the ‘052 Patent describes that:

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.**

Id. (emphasis added). Thus, the network transports data packets despite the data originating in a non-packet protocol (e.g., the AXI, OCP, DTL protocols).

121. This type of protocol conversion is not required in a bus because, as discussed above, a bus transmits data in the same protocol from which the data originated (e.g., AXI, OCP, DTL).

122. The figure below illustrates the difference between a multi-hop interconnect and traditional buses and point-to-point connections.



Ex. 2073 at 8

Each dot in the “network” above is a router or switch. Each router or switch is directly connected to its neighboring router or switch, but is indirectly connected to the remaining routers or switches through the “network.” Each router or switch also represents a “hop” (i.e., a router or switch that packets may traverse). “NoCs use packets to route data from the source to the destination [] via a network fabric that consists of switches (routers) [and] interconnection links (wires).” The bus shown above does not include routers or switches and does not transmit packets. Instead, the all functional units are directly connected to the bus via wires.

123. Thus, a POSITA would understand that a “hop” is a router or switch that a packet originating from an arbitrary protocol may traverse.

3. If “interconnect means” is subject to § 112, ¶ 6, the corresponding structure is a network

124. I understand that in the Institution Decision, the Board “construe[d] the corresponding structure for the ‘interconnect means’ to be a network, which is undisputed corresponding structure on this record.” Institution Decision at 18.

Assuming the term “interconnect means” is a means-plus-function limitation subject to § 112, ¶ 6, I agree: a POSITA would understand that the corresponding structure for “interconnect means” is a network.

125. In the Institution Decision, the Board stated that “Patent Owner has not proposed that we construe the corresponding structure for ‘interconnect means’ to be narrower than a network.” Institution Decision at 18. However, I understand that the Board did not expressly construe the term “network.” In the context of the ‘052 Patent, a POSITA would understand that a network, i.e., the corresponding structure for the interconnect means, is a “multi-hop interconnection between modules separated by one or more network nodes.” *Supra*, § VIII.A.2.

126. In the Institution Decision, the Board found that “the corresponding structure for the interconnect means should not include a mere bus, switch, and point-to-point wire.” Institution Decision at 18. I agree.

127. The ‘052 Patent specifically excludes a mere bus, switch, and point-to-point wire, alone, from being a network. *See* Ex. 1001 at 2:33-38 (“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. This is in contrast with the prevalent existing interconnects (i.e., buses) where modules are directly connected.”); *see also id.* at 1:50-54 (“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.”); *id.* at 4:46-52 (“The network on chip may include wires, bus, time-division multiplexing, switches, and/or routers within a network.”). A POSITA would not understand that the network “may include wires, bus, time-division multiplexing, switches, and/or routers” to mean that any one of those components, in isolation, can be a network. Instead, a POSITA would understand that the network could have those components, as long as they arranged to implement the multi-hop scheme and provided that they are separated by one or more network nodes.

128. Therefore, in the context of the ‘052 Patent, a POSITA would not understand the “interconnect means (N)” to be a mere “bus, switch, and point-to-point wire” as proposed by the Petitioner and Dr. Hassoun. Ex. 1003, ¶ 82. Instead, a POSITA would understand the “interconnect means (N)” to be a network, which a POSITA would understand has definite structure, and which is a “multi-hop interconnect separated by one or more network nodes.” For this reason, it is my

opinion that, if “interconnect means” is subject to § 112, ¶ 6, its structure is a “network” and a “network” is a “multi-hop interconnection separated by one or more network nodes.”

129. For these reasons, a POSITA understand that “interconnect means,” as used in the ‘052 patent, is a “multi-hop interconnection between modules separated by one or more network nodes.”

B. “mapping means” (claims 1 and 7)

130. I understand that Petitioner proposed the following construction for the term “mapping means,” which appears in claims 1-5 and 7”:

Term	Petitioner’s Construction
“mapping means (A) for 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 said at least one communication service identification”	Subject to Section 112, ¶ 6 Function: “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 said at least one communication service identification.” Structure: adapter unit/shell comprising a sequentializing unit and a desequentializing unit as disclosed in the ‘052 Patent at 5:58–6:3, 6:30–7:14, Figs. 1, 2A and 2B.

Petition at 11-12.

131. For purposes of this Declaration, I have been asked to apply this proposed construction in my analysis.

C. The Preamble of Claim 6

132. A POSITA would understand that the preamble of claim 6, reproduced below, is a limitation of the claim:

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:

Ex. 1001 at 10:7-18.

133. I understand that the preamble is often limiting if it provides the antecedent basis for terms in the body of the claim. The preamble of claim 6 does so in this instance. For example, the preamble provides antecedent basis for the following terms: (1) “a plurality of processing modules (M, S)”; (2) “at least one communication service”; (3) “specific communication properties”; and (4) “at least one communication service identification” recited in the body of claim 6.

134. A POSITA would understand that without the preamble being limiting, claim 6 would lack life, vitality, and meaning. Specifically, the body of the claim (i.e., limitation [6.d]) recites “mapping **the requested at least one communication service** based on **said specific communication properties**”

Ex. 1001 at 10:28-30 (emphasis added). The claimed “mapping” step requires “the requested at least one communication service based on said specific communication properties” which are defined and provided context in the preamble. Without the preamble describing that the “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 said specific communication properties,” a POSITA would not understand what “the requested at least one communication service” is. Therefore, the preamble is necessary to bring life, vitality, and meaning to claim 6.

135. I also understand that Petitioner in the related district court litigation agreed that the preamble of claim 6 should be construed as limiting. Ex. 2040 at 5-6. I also understand that the same construction was adopted by the district court in the TI case. Ex. 2022 at 58.

D. District Court Constructions

1. “connection properties” (claims 1, 3-7)

136. In the Qualcomm Litigation, I understand that the parties agreed that “connection properties” should be construed as “properties of a connection, including ordering (data transport in order), flow control (a remote buffer is reserved for a connection, and a data producer will be allowed to send data only when it is guaranteed that space is available for the produced data), throughput (a lower bound

on throughput is guaranteed), latency (upper bound for latency is guaranteed), the lossiness (dropping of data), transmission termination, transaction completion, data correctness, priority, or data delivery.” Ex. 2033 at Appendix 1; Ex. 2040 at 6.

137. I understand that the same construction was adopted by the district court in the TI case. See Ex. 2022 at 57.

2. “Integrated circuit” (claims 1-6)

138. In the Qualcomm Litigation, I understand that the parties agreed that “integrated circuit” should be construed as “interconnected circuitry on a chip.” Ex. 1016 at 13. I understand that the parties agreed to the same construction in the Samsung Litigation. Ex. 2033 at 5-6. I also understand that the TI Court construed “integrated circuit” as “a combination of interconnected circuit elements inseparably associated on or within a continuous substrate.” Ex. 2022 at 59.

3. “connection” (claims 1, and 3-7)

139. I am informed and understand that in the Qualcomm Litigation, the parties to that case proposed the following constructions for the term “connection.” Ex. 2040 at 6.

Term	Qualcomm’s Construction	Patent Owner’s Construction
“connection”	“one or more channels, each having a set of connection properties, between a first module and at least one second module, that are created before use and closed after use”	“one or more communication channels”

140. I understand that neither party is proposing a construction for the term “connection” in this proceeding. As such, I have been asked to apply the plain and ordinary meaning of this term to a POSITA at the time of the invention, considering the language of the claims, the specification, and prosecution history. However, if either of the parties’ district court construction were adopted in this proceeding, my opinions would not change.

IX. ANALYSIS OF THE ’052 PATENT

A. Ground 1: Rădulescu2004 (Claims 1-2 and 5-7)

141. In Ground 1, Dr. Hassoun opines that claims 1-2 and 5-7 of the ’052 Patent are anticipated by Rădulescu2004. Ex. 1003 at ¶¶ 120-153. For the reasons discussed below, I disagree.

142. At the outset, I note that the Board found that Petitioner did not establish a reasonable likelihood of success in connection with Ground 1 in the Institution Decision. Institution Decision at 30.

1. Rădulescu2004 Does Not Anticipate Independent Claim 6

143. In my opinion, Rădulescu2004 does not disclose all limitations of independent claim 6 and therefore does not anticipate.

i. Rădulescu2004 Does Not Disclose the Preamble of Claim 6

144. Dr. Hassoun opines that “Rădulescu2004 disclose the preamble of claim 6” and addresses the preamble in two pieces. Ex. 1003 at ¶ 122. For the reasons discussed below, I disagree.

a. 6.pre.1 “Method of communication service mapping in an integrated circuit having a plurality of processing modules (M, S)”

145. In discussing limitation [6.pre.1], Dr. Hassoun opines that Rădulescu2004 “teaches communication service mapping that is implemented on the chip, as I explain in further detail in my discussion of element [6.d].” Ex. 1003 at ¶ 122. As discussed below, I disagree with Dr. Hassoun with respect to limitation [6.d] and therefore disagree with his analysis for limitation [6.pre.1] for at least the same reasons. *Infra*, § VIII.A.1.i.c.

b. 6.pre.2 “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”

146. Limitation [6.pre.2] recites “at least one communication service identification” that “comprises at least one communication thread or at least one address range.” Dr. Hassoun does not opine that Rădulescu2004 discloses a communication service identification that comprises at least one communication thread. *See* Ex. 1003 at ¶¶ 124-127; Ex. 2068 at 90:3-6. Instead, Dr. Hassoun only opines that Rădulescu2004 discloses the claimed “address range.” *Id.* For the reasons discussed below, I disagree.

147. Dr. Hassoun opines that in Rădulescu2004, “the request is based on a communication service identification that is an address range as claimed” and that “Rădulescu2004 states that the requests from a master IP module include an address and/or memory within a slave module.” Ex. 1003 at ¶ 124. He then proceeds to identify several alleged examples of this. *Id.* I disagree; none of these examples discloses the claimed address range.

148. First, Dr. Hassoun quotes the following from § 1 of Rădulescu2004: “master IP modules issue request messages (e.g., read and write commands at an address, possibly carrying data) that are executed by the addressed slave modules.” Ex. 1003 at ¶ 124 (quoting Ex. 1005 at 1 § 1). Here, Rădulescu2004 is describing, at most, an address for a slave module. A POSITA would understand that an “address” is not the same thing as an “address range.” A POSITA would understand that a range has a lower bound and an upper bound. In other words, a POSITA would understand that the address range in a communication service identification specifies a continuous block or series of addresses. When a communication service identification comprises an address range, the communication service request identifies one or more particular locations with the slave functional block for receiving and performing the request. For example, if the address range is “0000-7FFF,” the first portion (0000) specifies the lower bound while the second portion (7FFF) specifies the upper bound of the range. Ex.1001 at 8:12-15.

149. The distinction between an address and an address range is particularly relevant to a narrowcast connection, which a POSITA would understand is a connection between one master and multiple slaves. Ex. 2074 at 12:53-60; Ex. 1005 at 2 § 2. When the communication service request of an eventual narrowcast connection specifies an address range as part of the

communication service identification, a POSITA would understand that (1) the slave identified by the address range is the target of the request and (2) there are one or more modules within the slave's address range that perform the request. This is in contrast to the "address" disclosed by Rădulescu2004 which a POSITA would understand only identifies one slave or one location within the slave. In other words, Rădulescu2004's "address" is incapable of specifying more than one location within the targeted slave to perform the request. To specify more than one location within the slave, Rădulescu2004's system requires the master to issue multiple requests, each request specifying a particular address, even though the address in each of the requests may be directed to the same slave.

150. The disclosure in Rădulescu2004 relied on by Dr. Hassoun as purportedly disclosing an "address range" actually highlights why Rădulescu2004 does not in fact disclose an address range. Ex. 1005 at 1 § 1 ("master IP modules issue request messages [plural] (e.g., read and write commands at an address, possibly carrying data) that are executed the by the addressed slave modules [plural]."). Based on this disclosure, 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*.

151. Second, Dr. Hassoun asserts that Rădulescu2004 discloses the claimed “address range” because Rădulescu2004 discloses “Masters initiate transactions by issuing request, which can be further split in commands, and write data ... slaves receive and execute each transaction.” Ex. 1003 at ¶ 124 (citing Ex. 1005 at 2 § 2). Nothing in this portion of Rădulescu2004’s disclosure suggests a communication service identification comprising an address range. As discussed above, the “address range” is used to specify one or more locations within the targeted slave’s address range that perform the issued request. This portion of Rădulescu2004 does not disclose this concept.

152. Third, Dr. Hassoun asserts that Rădulescu2004 discloses the claimed “address range” because Rădulescu2004 discloses that “examples of commands are read and write.” Ex. 1003 at ¶ 124 (citing Ex. 1005 at 2 § 2). Nothing in this portion of Rădulescu2004’s disclosure suggests a communication service identification comprising an address range. Specifically, Rădulescu2004 does not disclose that there are one or more locations within the targeted slave’s address range that perform the “commands.”

153. Fourth, Dr. Hassoun asserts that Rădulescu2004 discloses the claimed “address range” because Rădulescu2004 discloses that “communication is performed using read/write transactions.” Nothing in this portion of Rădulescu2004’s disclosure suggests a communication service identification

comprising an address range. Specifically, Rădulescu2004 does not disclose that there are one or more locations within the targeted slave's address range that perform the "read/write transactions."

154. Fifth, Dr. Hassoun relies on Rădulescu2004's "Figs. 3-6 and accompanying description" as purportedly disclosing the claimed "address range." I disagree. I note that Dr. Hassoun does not specifically identify any portions of these figures or their description that purportedly disclose an address range. I have reviewed these figures and corresponding description and have concluded that none of them disclose or suggest an address range. Specifically, the description accompanying Figure 3 makes clear that the "address range assigned to a slave is configurable in the narrowcast module." Ex. 1005 at 4 § 4.2. As explained below, a POSITA would understand that an address range is not communicated in a request if it "is configurable in the narrowcast module." The description accompanying Figure 4 is directed to a "multi-connection shell," and does not mention an "address range." *Id.* Similarly, the description accompanying Figures 5 and 6 is directed to "master and slave shells that implement a simplified version of a protocol such as AXI" and does not mention an "address range." *Id.* at 5 § 4.2.

155. Nowhere in Rădulescu2004 is it disclosed that the communication service request issued by the master that is based on a communication service identification comprising an address range. Specifically, Rădulescu2004 does not

disclose that there are one or more locations within the targeted slave's address range that perform the issued request.

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 Rădulescu2004 discloses an address (“addr”) and not an address range.

157. Dr. Hassoun's opinion with respect to Ground 2 also highlights how the “address” of Rădulescu2004 and the claimed “address range” are different. Specifically, Dr. Hassoun opines that OCP (Ex. 1006) purportedly discloses the claimed address range because it discloses the “MAddrSpace” signal. Ex. 1003 at ¶ 172 (citing Ex. 1006 at 38). Dr. Hassoun emphasizes OCP's disclosure that the

“MAddrSpace” signal “is used to indicate the address region of a transfer,” while the “MAddr specifies the slave-dependent address of the resource targeted by the current transfer.” *Id.* (citing Ex. 1006 at 24, 26). I understand that Dr. Hassoun’s opinion with respect to Ground 2 concedes that OCP’s “MAddr” signal is different than the “MAddrSpace” signal; therefore, the “address” of Rădulescu must be different than the “address range” claimed in the ‘052 Patent.

158. Dr. Hassoun also states that “Rădulescu2004 also explains that each slave identified by an address has an ‘address range’ associated with it, that address range is ‘configurable,’ and that the slave that the transaction is destined for is ‘selected based on the address.’” Ex. 1003 at ¶ 124 (quoting Ex. 1005 at § 4.2). This statement in Rădulescu2004 does not disclose the communication service request including a communication service identification comprising an “address range.”

159. Rădulescu2004 states that “[n]arrowcast 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.***” Ex. 1005 at 4 § 4.2. Rădulescu2004 also explains that “the slave for which the transaction is destined is selected ***based on the address*** (CONN block)” and that “[t]he address range assigned to a slave is configurable in the narrowcast module.” Ex. 1005 at 4 § 4.2. Dr. Hassoun conflates two separate concepts in Rădulescu2004: (1) that a

request includes an address and (2) that a slave can be assigned an address range.

Ex. 1005 at 1-4, FIG. 5. However, even if these two concepts are put together, Rădulescu2004 still does not disclose that the “communication service identification comprises ... at least one address range” because the master module (M) “requests at least one communication service.” 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.

160. I understand that the Board preliminarily found that “Rădulescu2004’s disclosures that its communication service identification contains an address for the ‘slave’ and its ‘slave’ has an address range do not by themselves teach or suggest that Rădulescu2004’s communication service identification contains an address range.” Institution Decision at 25. For the reasons discussed above, I agree.

161. For the reasons provided above, Rădulescu2004 does not disclose a “communication service identification, wherein said at least one communication service identification comprises at least one communication thread or at least one address range,” as required by the preamble of claim 6.

ii. 6.d “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 said at least one communication service identification.”

162. Rădulescu2004 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 said at least one communication service identification,” as recited in claim 6 of the ‘052 Patent because Rădulescu2004 does not disclose the claimed “communication service identification.” *Supra*, § IX.A.1.i.b.

163. Even if Rădulescu2004 discloses the claimed communication service identification (which it does not), Rădulescu2004 still does not disclose the claimed “mapping” step. Dr. Hassoun opines that Rădulescu2004 discloses the “mapping” step because “each transaction is executed by a single slave selected based on the address provided in the transaction.” Ex. 1003 at ¶ 143 (quoting Ex. 1005 at 4 § 4.2). I disagree. A POSITA would also understand that merely selecting a slave “based on the address provided in the transaction,” as described in Rădulescu2004 does not satisfy the “mapping” step claimed in the ‘052 Patent.

164. The “mapping” step requires: “**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 said at least one**

communication service identification.” Ex. 1001 at 10:28-32. In particular, Rădulescu2004 lacks the claimed “mapping ... to a connection based on a set of connection properties.” Rădulescu2004 discloses that, in a narrowcast connection (i.e., a connection between one master and several slaves), the target slave(s) are selected based on their respective addresses. Rădulescu2004 does not disclose that the communication service is mapped to a connection according to the communication service identification (which Rădulescu2004 fails to disclose, *supra* § IX.1.i.b) and “**based on a set of connection properties.**”

165. I reproduce Rădulescu2004’s relevant disclosure below:

In Figure 3, we show an example of a narrowcast shell. Narrowcast 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 [23]. Narrowcast connections provide a simple, low-cost solution for a single shared address space mapped on multiple memories. It implements the splitting/merging of data going to/coming from these memories.

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. In-order delivery per slave of request messages is already provided by the point-to-point connections.

Ex. 1005 at 4 § 4.2; Ex. 1003 at ¶ 140.

166. There is no mention of selecting a specific connection based on connection properties assigned to that connection. The specification provides that, when the “mapping” step includes a communication service identification comprising an “address range,” the “network interface can identify each stream, and **map it to the right connection having connection properties.**” Ex. 1001 at 8:33-36 (emphasis added). Similarly, when the “mapping” step includes a communication service identification comprising a thread, “such a thread and the associated traffic class can be mapped to a connection based on **specific corresponding connection properties.**” *Id.* at 7:44-48 (emphasis added). As confirmed by the specification, a POSITA would understand that the claimed “mapping” step requires mapping the requested communication service to a specific connection, the specific connection being based on specific connection properties.

167. In contrast, the only “connection property” disclosed in Rădulescu2004 is ordering, and Rădulescu2004 makes clear that the ordering of requests and responses is not a property of the “connection,” but rather a property of network implementation. With respect to ordering of responses, Rădulescu2004 discloses that “the narrowcast must [] keep a history of connection identifiers of the transactions including responses ... and the length of these responses.” Ex. 1005 at 4 § 4.2. Thus, ordering of responses is controlled by the network and the

network interface and is not a property of the connection. Similarly, with respect to ordering of requests, Rădulescu2004 discloses that “[i]n-order delivery per slave of request messages is already provided by the point-to-point connections.” A POSITA would understand that this means that the network itself implements in-order transmission of response messages (i.e., the first message is transmitted first and the last message is transmitted last). Thus, ordering of requests is controlled by the network and network interface and is not a property of the connection.

168. Dr. Hassoun opines that the specification and the prosecution history limit the “mapping” to merely “mapp[ing] to [] connections corresponding to those slaves based on their address ranges.” Ex. 1003 at ¶¶ 141-142. I disagree.

169. Dr. Hassoun’s opinion is contrary to how the “mapping” is described in the ‘052 Patent. As one example of the “mapping,” the ‘052 Patent provides:

[T]wo communication threads identifying two traffic classes, namely one to a fast memory, and the other on a slow memory, may be considered. **The thread identifying the fast traffic class can be mapped to a connection configured to guarantee high bandwidth, while the slow traffic class can be mapped to a connection with a low guaranteed bandwidth.**

Ex. 1001 at 7:49-55 (emphasis added).¹ A POSITA would understand that “mapping ... to a connection based on a set of connection properties” requires a

¹ Rădulescu2004 does not disclose a “thread;” however, the ‘052 Patent’s exemplary “mapping” step involving threads is relevant to show how the claimed “mapping” is different from Rădulescu2004’s disclosure.

consideration of the connection properties associated with the connection to effectuate the “communication properties” associated with the “communication service.” In other words, “mapping” requires more than merely selecting any connection, which is what Rădulescu2004 discloses. This concept of “mapping” to a particular connection having a set of specific connection properties is explicitly required by the ‘052 Patent and forms the basis of the invention:

[T]he system comprises **a mapping facility to map a range of addresses or a thread information to connection data.** The connection data comprises specifications for data transport such as best effort or guaranteed throughput, the maximum allowed jitter, the maximum allowed latency, flow control Y/N, or ordering Y/N. **If a message is sent to a particular address range in the system, or belongs to a particular thread, a connection is assigned to assure that the message is transmitted via a particular set of transmission requirements.**

Ex. 1001 at 8:64-9:5 (emphasis added).

170. Rădulescu2004 does not disclose limitation [6.d] because Rădulescu2004 does not consider the “connection data” when selecting the connection for the narrowcast transmission. In other words, Rădulescu2004 does not disclose mapping a communication service to a particular connection “to assure that the message is transmitted via a particular set of transmission requirements.”

Ex. 1001 at 8:64-9:5. Instead, Rădulescu2004 merely discloses selecting the slave “based on the address provided in the transaction” without regard to the “particular set of transmission requirements.”

171. Thus, Rădulescu2004 does not disclose “mapping ... to a connection based on a set of connection properties,” as claimed in claim 6 of the ‘052 Patent.

2. Rădulescu2004 Does Not Anticipate Claims 1-2, 5, and 7

172. I understand that Dr. Hassoun opines that Rădulescu2004 discloses the limitations of claims 1-2, 5 and 7 for the same reasons discussed for independent claim 6. Ex. 1003 at ¶¶ 144-153. I disagree with Dr. Hassoun’s conclusion for claims 1-2, 5, and 7 for the same reasons discussed above for claim 6. *Supra* § IX.A.1. Specifically, Rădulescu2004 fails to disclose “at least one communication service identification compris[ing] at least one communication thread or at least one address range;” and “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 said at least one communication service identification.”

B. Ground 2: Purported Obviousness over Rădulescu2004 in view of OCP (Claims 1-7)

173. For Ground 2, Dr. Hassoun opines that claims 1-7 of the ‘052 Patent are obvious over Rădulescu2004 and OCP. Ex. 1003 at ¶¶ 154-182. For the reasons discussed below, I disagree.

1. A POSITA Would not Combine Rădulescu2004 and OCP

174. The ‘052 Patent is directed to advancements over modules communicating in a bus-based protocol, such as the OCP protocol. Specifically,

the ‘052 Patent recognized the deficiencies of conventional SoCs that used the OCP protocol, while also recognizing the wide adoption of the OCP protocol and strived to improve the interconnect that connects the functional blocks communicating in the OCP protocol. Ex. 1001 at 2:43-45 (“Network on chip are a platform-based design with an aim to reduce the cost of system design through re-use of applications and architectures); *see also id.* at 4:15-18 (“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 ‘052 Patent is directed to improvements of effectuating the communication between functional blocks communicating in existing protocols, such as, for example, the OCP protocol.

175. Dr. Hassoun opines that “OCP teaches communication service requests based on communication service identifiers that include both a thread and an address range” and that a POSITA would be motivated to incorporate each individually into Rădulescu2004’s system. Ex. 1003 at ¶¶ 172-173. I disagree. A POSITA would not be motivated to use either threads or address ranges in Rădulescu2004’s system.

176. With respect to address range, Dr. Hassoun’s appear to be based on an incorrect assumption that OCP’s MAddrSpace signal (which he alleges is analogous to the claimed address range) is a necessary signal sent in every OCP

request. Ex. 1003 at ¶ 172. That assumption would be incorrect. OCP provides that “[e]ach OCP interface on the core must ... [c]ontain at least the basic OCP signals.” Ex. 1006 at 15. The basic signals included in each communication include Clk, MAddr, MCmd, MData, SCmdAccept, SData, and SResp. *Id.* at 24. The MAddr signal “specifies the slave-dependent address of the resource targeted by the current transfer.” *Id.* In contrast, the MAddrSpace signal is one of many “optional dataflow signals [that] are grouped into simple and complex extensions.” *Id.* at 24, 26. MAddrSpace is “used to indicate the address region of a transfer.” *Id.* at 26. Based on OCP’s disclosure, a POSITA would understand that the “MAddrSpace” signal (i.e., the signal that Dr. Hassoun opines is analogous to the claimed “address range”) is an optional extension of the OCP protocol and is not necessarily required.

177. Dr. Hassoun fails to articulate why a POSITA would be motivated to use the optional MAddrSpace signal extension in Rădulescu2004’s system. Ex. 1003 at ¶¶ 172-173. Dr. Hassoun merely opines that a “POSA would have understood that MAddr and MaddrSpace identify a system target (*i.e.*, second processing module) and/or a memory region within the system target” (*id.* at ¶ 172) and concludes that, for this reason, “a POSA would have been motivated to implement the OCP protocol, including its ... address range teachings, in the system of Rădulescu2004” (*id.* at ¶ 173). I disagree. As I explained above,

Rădulescu2004 discloses a communication service identifier in the form of an address (not the claimed “address range”) to select “the slave for which the transaction is destined.” *Supra* § IX.A.1.i.b; Ex. 1005 at 4 § 4.2. A POSITA would understand that OCP’s MAddr signal is analogous to the address disclosed in Rădulescu2004.

178. Rădulescu2004 also provides that “narrowcast connections provide a simple, low-cost solution for a single shared address space mapped on multiple memories.” Ex. 1005 at 3 § 4. A POSITA would understand that Rădulescu2004’s narrowcast connection already provides a means for “identify[ing] a ... memory region within the system target.” Ex. 1003 at ¶ 172. Therefore, a POSITA would know that the optional MAddrSpace extension in OCP is not required in Rădulescu2004’s system to identify a memory region. A POSITA would find no benefit in using the optional MAddrSpace extension in Rădulescu2004’s system, therefore, a POSITA would not be motivated to combine OCP and Rădulescu2004 in the way Dr. Hassoun proposes.

179. Dr. Hassoun fails to explain why a POSITA would opt to use MAddrSpace signal when OCP’s standard MAddr signal (i.e., an “address” present in Rădulescu2004) already provides Dr. Hassoun’s alleged benefits of using the MAddrSpace signal, i.e., that it would identify a system target and/or a memory region within the system target.

180. For these reasons, I disagree with Dr. Hassoun that a POSITA would be motivated to use OCP's optional MAddrSpace extension in Rădulescu2004's system.

181. With respect to threads, Dr. Hassoun opines that OCP teaches threads and, because of this teaching, "a POSA would have been motivated to implement the OCP protocol, including its threads ... teachings, in the system of Rădulescu2004." Ex. 1003 at ¶¶ 169-171, 173. I disagree. Dr. Hassoun fails to opine as to what would motivate a POSITA to implement threads into Rădulescu2004's system besides the mere fact that OCP teaches threads. *Id.* Instead, Dr. Hassoun merely opines that "OCP's teaching regarding thread properties are consistent with the '052 patent." *Id.* at ¶ 171.

182. I disagree with Dr. Hassoun's conclusions because they appear to be based on the incorrect assumption that OCP's MThreadID signal (which he alleges is analogous to the claimed communication thread) is a necessary signal sent in every OCP request. Ex. 1003 at ¶ 172. That assumption would be incorrect. OCP provides that "[e]ach OCP interface on the core must ... [c]ontain at least the basic OCP signals." Ex. 1006 at 15. The basic signals included in each communication include Clk, MAddr, MCmd, MData, SCmdAccept, SData, and SResp. *Id.* at 24. The MThreadID signal is a "request thread identifier." *Id.* The MThreadID signal is one of many "optional dataflow signals [that] are grouped into simple and

complex extensions.” *Id.* at 24, 27. Specifically, MThreadID is a complex extension to the OCP protocol. *Id.* at 27. MThreadID is “used to indicate the address region of a transfer.” *Id.* at 26. Based on OCP’s disclosure, a POSITA would understand that the “MThreadID” signal (i.e., the signal that Dr. Hassoun opines is analogous to the claimed “communication thread”) is an optional extension of the OCP protocol and is not necessarily required.

183. There is no description or teaching in Rădulescu2004 regarding threads. Dr. Hassoun explains that OCP’s thread-based requests include communication properties such as ordering and priority, but he fails to explain why a POSITA would be motivated to implement such thread-based communication properties into Rădulescu2004’s system. Ex. 1003 at ¶ 169. Instead, Dr. Hassoun merely opines that the ‘052 Patent teaches similar communication properties. *Id.* at ¶ 170.

184. A POSITA would not be motivated to implement OCP’s thread-based teachings into Rădulescu2004 because Rădulescu2004 already teaches communication properties. Ex. 1005 at 1 § 1 (“We provide transport-layer services by defining connection (e.g., point-to-point or multicast) configured for specific properties (e.g., throughput, ordering); *id.* at 2 § 2 (“different properties can be attached to the request and response parts of a connection.”). Specifically, Dr. Hassoun opines that OCP’s threads provide ordering. Ex. 1003 at ¶¶ 169-171.

However, Rădulescu2004 provides that ordering is already implemented in its system without the use of threads. Ex. 1005 at 4 § 4.2 (“To provide in order response delivery, the narrowcast module 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. In-order delivery per slave of request messages is already provided by the point-to-point connections.”). Moreover, a POSITA would not be motivated to use OCP’s threads in Rădulescu2004 as a form of communication service identification because Rădulescu2004 teaches that “the slave for which the transaction is destined is selected based on the address.” *Id.* at 4 § 4.2. In other words, a POSITA would not recognize OCP’s disclosure of threads to be a form of communication service identification missing from Rădulescu2004 because neither OCP nor Rădulescu2004 provide any reason for using OCP’s threads as a form of communication service identification.

185. Dr. Hassoun also opines that “OCP teaches communication service requests based on communication service identifiers that include both a thread and an address range.” Ex. 1003 at ¶ 172. Dr. Hassoun relies on Table 12 of OCP for his opinion. *Id.* Table 12 is reproduced below.

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

Ex. 1006 at 38.

I note that the optional simple and complex extensions described above, namely, MAddrSpace and MThreadID remain optional signals despite their condition being “always.” Table 12 of OCP merely states that fields are grouped together (i.e., the “request” group, the “response” group, and the “datahandshake” group). Ex. 1006 at 38. MAddrSpace and MThreadID are merely examples of signals that could be in the “request” group. *Id.*

186. For these reasons, I disagree with Dr. Hassoun that a POSITA would be motivated to use OCP’s threads in Rădulescu2004’s system.

187. I also disagree that a POSITA would generally be motivated to combine the teachings of Radulecsu2004 with OCP. I also note that OCP is a connectionless communication protocol. While I understand that OCP uses the term “connection,” that term is a misnomer and does not exemplify the type of “connection” disclosed in Rădulescu2004 and the ‘052 Patent. *See e.g.*, Ex. 1006 at 120 (“The components of a chip are joined together using connections. A connection statement instantiates a *bundle*, which is a group of signals.”) (emphasis in original). This is because network-based connections (such as those disclosed in Rădulescu2004 and the ‘052 Patent) transmit data in the form of packets or flits, while the “connection” disclosed in OCP transmits data in the form of a transaction (i.e., via the OCP protocol). *See supra* § VIII.A. OCP also provides a “networking analogy” which emphasizes OCP’s lack of packets. Ex. 1006 at 170 (“Using a networking analogy, the thread ID is a level-2 (**data link layer**) concept, whereas the connection ID more like a level-3 (**transport/session layer**) concept.”) (emphasis added). A POSITA would understand that packets are not transmitted in the data link, transport, or session layers of the OSI-model; instead, packets are transmitted in the network layer of the OSI-model. The ‘052 Patent specifically discloses a packet-based communication protocol. In particular, the ‘052 Patent provides:

Summarizing it can be said, that the above described mapping is relevant for communication protocols which **do not include the notion of connection, but are implemented with an interconnect supporting connections** The advantage is that **it enables the user of connections for protocols without the notion of connection,** allowing the configuration of differential services for different traffic classes. Backward compatibility of on-chip networks is achieved, by means of the adaptor units/shells **to enable existing protocols,** such as DTL, AXI, and **OCP,** to use differential services **as offered by a connection based interconnect, such as an on-chip network.**

Ex. 1001 at 8:47-58 (emphasis added). A POSITA would similarly understand that the “connection” described in the ‘052 Patent and Rădulescu2004 is different than the connection described by OCP.

188. Because OCP and Rădulescu2004 are fundamentally different (except with respect to OCP’s functional blocks), I disagree with Dr. Hassoun’s opinions that rely on OCP’s connections being synonymous with the ‘052 Patent’s connections. Specifically, Dr. Hassoun opines that “both Rădulescu2004 and OCP use connection id signals in their interfaces to map a particular message to a particular connection.” Ex. 1003 at ¶ 158. While I disagree that Rădulescu2004 discloses “mapping” (*supra*, § IX.A.2) even if it did, the mapping would not be analogous to the mapping disclosed in OCP. Specifically, OCP instructs:

Use connections whenever information about a request must be sent end-to-end from master to slave. Any bridges in the path between the end-to-end partners preserve the connection ID, even as thread IDs are re-assigned on each OCP interface in the path. The MConnID field transfers the connection ID during the request phase. Since this

establishes the mapping onto a thread ID, the other phases do not require a connection ID but are unambiguous with only a thread ID.

Ex. 1006 at 47. There are several reasons why this statement is not analogous to the “mapping” disclosed in the ‘052 Patent and the alleged “mapping” disclosed in Rădulescu2004.

189. First, the connection ID is mapped onto the thread ID as opposed to the request being mapped onto a connection having a particular set of connection properties according to the communication service identification (e.g., a thread ID), as required by the ‘052 Patent. Ex. 1001 at 10:28-32. In other words, OCP teaches that the connection ID is incorporated into the thread ID throughout the transaction. This is opposite to the disclosure of the ‘052 Patent. This disclosure also emphasizes that the “connection” disclosed by OCP is not the same as the “connection” disclosed by Rădulescu2004 and the ‘052 Patent because, at least in the ‘052 Patent, the connection is meant to “decouple computation for communication,” (*id.* at 1:42-43), i.e., by implementing connections, the properties of a service can be abstracted from the functional block initiating the request. This is in contrast to OCP which discloses that the connection ID and the thread ID are two signals originating from the initiator. Ex. 1006 at 38, Table 12.

190. Second, the “mapping” in OCP occurs at the OCP interfaces. *See* Ex. 1006 at 47. As discussed above, any disclosure related to OCP interfaces is

irrelevant to this analysis. To reiterate, this is because Rădulescu2004 makes clear that the only relevant disclosure in OCP is the functional blocks, which could be used with Rădulescu2004's network interfaces and network.

191. Moreover, I also disagree with Dr. Hassoun's suggestion that "any on-chip interconnect can be interfaced to the OCP." Ex. 1003 at ¶ 161 (quoting Ex. 1006 at 11). Version 1.0 of OCP (i.e., Ex. 1006) published prior to the commercialization of network-on-chip technology. *Compare* Ex. 1006 at 4 (copyrighted in 2001) *with* Ex. 2072 at 9 (first NoC commercially available in 2006). Therefore, at the time of OCP's publication, it may be true that any on-chip interconnect could be interfaced to OCP. However, following the development and commercialization of NoC, that was not true until the NoC disclosed in Rădulescu2004 and an advancement of the same NoC described in the '052 Patent. The objective of the '052 Patent was to provide compatibility between a NoC and the OCP protocol. In other words, a POSITA would understand that the disclosure in OCP that "any on-chip interconnect can be interfaced to the OCP" was incorrect at the time of the invention of the '052 patent because NoCs were not commercially available when OCP was first published.

192. For at least these reasons, a POSITA would not be motivated to combine all of OCP's teachings with Rădulescu 2004. Instead, a POSITA would look solely at OCP's description of functional block communication and would not

consider how the OCP interfaces and interconnects operate in conjunction with the functional units.

2. Rădulescu2004 in View of OCP Does Not Render Obvious Independent Claim 6

193. Dr. Hassoun opines that Rădulescu2004 “teaches each element of the claim,” including “a communication service identifier in the form of an address range.” Ex. 1003 at ¶ 163. For the reasons articulated above, I disagree that Rădulescu2004 discloses each limitation of claim 6 of the ‘052 Patent. *Supra* § IX.A.1.

194. Dr. Hassoun further opines that OCP discloses a thread as the claimed communication service identifier. Ex. 1003 at ¶ 163. I disagree that the combination of Rădulescu2004 and OCP discloses each element of the claim and I further disagree that Rădulescu2004 in view of OCP renders claim 6 obvious.

- i. **6.pre.2 “wherein at least one first of said processing modules (M) request 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)”**

195. Limitation [6.pre.2] recites “wherein said at least one communication service identification comprises at least one communication thread or at least one address range.” Dr. Hassoun opines that “Rădulescu2004 teaches this part of the claim.” Ex. 1003 at ¶ 163. For the reasons provided above, I disagree. *Supra* § IX.A.1.i.b. Dr. Hassoun also opines that “a communication service identifier in the form of a communication thread would have been obvious in view of OCP.” Ex. 1003 at ¶ 163. I disagree. Specifically, a POSITA would not be motivated to combine Rădulescu2004 with OCP’s disclosure of threads and address ranges because the purported benefits of such combination are already realized in Rădulescu2004 without threads or address ranges. *Supra* § IX.B.1. Because a POSITA would not be motivated to make such a combination, Rădulescu2004 in view of OCP does not render obvious limitation [6.pre.2].

ii. 6.d “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 said at least one communication service identification.”

196. Limitation [6.d] recites “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 said at least one communication service identification.” Dr. Hassoun opines that Rădulescu2004 in view of OCP renders obvious limitation [6.d]. Ex. 1003 at ¶¶ 174-177. I disagree.

197. As I explained above, Rădulescu2004 fails to disclose the claimed “communication service identification.” *Supra* § IX.A.1.i.b. I also explained that a POSITA would not be motivated to incorporate OCP’s teaching of MAddrSpace and threads into Rădulescu2004. *Supra* § IX.B.1. For these reasons, the combination of Rădulescu2004 and OCP do not render obvious limitation [6.d].

198. Rădulescu2004 also fails to disclose the claimed “mapping ... to a connection based on a set of connection properties.” *Supra* § IX.A.1.ii. OCP fails to cure the deficiencies of Rădulescu2004 with respect to the claimed “mapping ... to a connection based on a set of connection properties” because, as explained above, a POSITA would not combine OCP’s teachings regarding “mapping” with Rădulescu2004. *Supra* § IX.B.1. Thus, Rădulescu2004 in view of OCP cannot

render obvious “mapping ... to a connection based on a set of connection properties.”

199. Dr. Hassoun opines that it would have been obvious to include “an OCP adapter shell in the network interface of Rădulescu2004.” Ex. 1003 at ¶ 174 (citing Ex. 1005 at FIG. 1). Dr. Hassoun provides no explanation for why a POSITA would find this to be obvious. I note that Rădulescu2004 does not disclose an OCP adapter shell in the network interface. Ex. 1005 at FIG. 1 (reproduced below).




Figure 1. NI kernel and shells

Ex. 1005 at FIG. 1.

All of Dr. Hassoun’s opinions with respect to the combination of Rădulescu2004 and OCP are based on a POSITA finding it obvious to include an OCP adapter in FIG. 1 of Rădulescu2004. *See* Ex. 1003 at ¶¶ 174-177. However, neither Rădulescu2004 nor OCP disclose an OCP adapter shell. *See generally* Ex. 1005;

Ex. 1006. Therefore, a POSITA would have to create and design an OCP adapter shell in order to include it in Rădulescu2004's network interface. In my opinion, this would not be obvious to a POSITA.

200. Based on his assumption that a POSITA would include an OCP adapter shell in Rădulescu2004's network interface, Dr. Hassoun further opines that it would be obvious "to use the MAddr / MAddrSpace fields of OCP in the request message ... and communicate it via the msg_req signal ... in order to implement the Rădulescu2004 disclosure of configuration of address range assigned to a slave as part of a narrowcast connection. Ex. 1003 at ¶ 175 (citing Ex. 1005 at FIGs. 3, 7). While this sentence is difficult to follow, I understand Dr. Hassoun is opining that it would have been obvious for a POSITA to use the MAddrSpace signals of OCP in Rădulescu2004's request message. However, Dr. Hassoun fails to identify any motivation for making this combination. A POSITA would not be motivated to make this combination given Rădulescu2004's disclosure that the target slave is identified "based on the address provided in the transaction." Indeed, a POSITA would understand that identifying the targeted slave "based on the address provided in the transaction" is a permissible way of identifying the targeted slave. In other words, the purpose of the "address" in Rădulescu2004 is to identify the targeted slave module of the request. Dr. Hassoun provides no reasoning behind incorporating the "MAddrSpace" signals of OCP

into Rădulescu2004. A POSITA would understand that the “MAddrSpace” signals would identify a slave; however, a POSITA would also understand that the “address” of Rădulescu2004 also identifies a slave. There are no deficiencies in Rădulescu2004 that render Dr. Hassoun’s proposed combination obvious.

201. Dr. Hassoun further opines that “mapping requests to connections according to a communication service identification that includes a Thread ID ... was something that a POSA would have been motivated to do with a reasonable expectation of success, in view of the teachings of Rădulescu2004 and OCP.” Ex. 1003 at ¶ 176. I disagree.

202. Neither Rădulescu2004 nor OCP disclose the claimed “mapping.” *Supra* §§ IX.A.1.ii, IX.B.1. Dr. Hassoun does not opine that OCP discloses the claimed “mapping.” *See* Ex. 1003 at ¶¶ 174-177. Dr. Hassoun also fails to opine as to whether “mapping ... to a connection based on a set of connection properties” is disclosed by either Rădulescu2004 or OCP. *Id. Supra* § IX.A.1.ii. Instead, Dr. Hassoun merely opines “mapping requests to connections according to a communication service identification” would have been obvious. Ex. 1003 at ¶ 176. As discussed above, this is not what the claim requires. *Supra* § IX.A.1.ii. Instead, the claim requires “mapping ... to a connection based on a set of connection properties.” Neither Rădulescu2004 nor OCP disclose “mapping ... to a connection based on a set of connection properties.” For these additional reasons,

Dr. Hassoun’s opinions with respect to “mapping ... to a connection based on a set of connection properties” is deficient.

203. Dr. Hassoun further opines that “OCP teaches using ‘connections’ for the transfer of requests with specified communication properties over an interconnect—just like Rădulescu2004.” Ex. 1003 at ¶ 176. As I explained above, *supra* § IX.B.1, the “connections” described in OCP are different than the “connections” described in Rădulescu2004. A POSITA would understand that these connections are not comparable and would not be motivated to take the teachings of OCP with respect to “connections” and apply them to Rădulescu2004.

204. Dr. Hassoun further opines that a POSITA would have been motivated “to map **both** the OCP address signals and OCP thread identifier to connections of Rădulescu2004, because the functionality of these features are not in conflict with each other, serve complimentary purposes.” Ex. 1003 at ¶ 177. Again, I note that Dr. Hassoun fails to opine whether a POSITA would have found it obvious to perform the step of “mapping ... to a connection based on a set of connection properties” as required by claim 6 of the ‘052 Patent. In fact, Dr. Hassoun fails to discuss connection properties at all with respect to limitation [6.d]. *See* Ex. 1003 at ¶¶ 174-177.

205. Dr. Hassoun further opines that “by mapping address signals to Rădulescu2004’s connections, the NoC can identify which particular slave module

is being targeted by a narrowcast communication service request.” Ex. 1003 at ¶ 177. To the extent that Dr. Hassoun is opining that it would be obvious to use the MAddrSpace signal disclosed in OCP in Rădulescu2004’s request, I disagree for the reasons provided above. *Supra* § IX.B.1. Specifically, a POSITA would find no reason to incorporate the MAddrSpace signal disclosed in OCP into Rădulescu2004’s NoC because Rădulescu2004 does not require it to achieve its purpose of identifying the target slave “based on the address provided in the transaction.” Ex. 1005 at 4, § 4.2. In other words, Dr. Hassoun’s purported motivation for this combination is it will allow the NoC to “identify which particular slave module is being targeted by a narrowcast communication service request.” Ex. 1003 at ¶ 177. A POSITA would not be motivated to make this combination because Rădulescu2004 already discloses that “[w]ithin a narrowcast connection, the slave for which the transaction is destined is selected based on the address.” A POSITA would realize that there is no benefit gained from this combination.

206. Dr. Hassoun also opines that “by mapping OCP threads to Rădulescu2004’s connections, the NoC can perform thread-based ordering and priority can be communicated through the NoC ‘between a system initiator and a system target,’ and implement Rădulescu2004’s multiple connections to a port in order to allow differentiated traffic classes through the connid signals.” Ex. 1003 at

¶ 177 (citing Ex. 1006 at 180-181; Ex. 1005 at § 4.1). I disagree that a POSITA would be motivated to make this combination. In particular, Rădulescu2004 already discloses a means for providing ordering. Ex. 1005 at 4 § 4.2 (“To provide in-order response delivery, the narrowcast must also keep a history of connection identifiers of the transactions including responses ... In-order delivery per slave of request messages is already provided by the point-to-point connections.”). Similarly, as Dr. Hassoun admits, Rădulescu2004 already provides differentiated traffic classes. Ex. 1003 at ¶ 166 (“Rădulescu2004 also teaches that one of the benefits of its NoC system is that it allows ‘differentiated services’ and ‘differentiated traffic classes’”); Ex. 1005 at 3 § 4.1 (“A port can have multiple connections to allow differentiated traffic classes”).

207. Dr. Hassoun also opines that a POSITA would be motivated to include Thread IDs in this proposed combination because Thread IDs provide “priority” information that can be communicated through the NoC. Ex. 1004 at ¶ 177 (citing Ex. 1006 at 180-181). I disagree that a POSITA would be motivated to make this combination. In particular, Rădulescu2004 already discloses a means for providing priority information. Ex. 1005 at 2 § 2 (“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**.”). A POSITA would

understand that providing guaranteed message delivery and timing guarantees is equivalent to providing priority information.

208. Therefore, the alleged benefits of Dr. Hassoun’s proposed combinations (both with respect to incorporating the MAddrSpace signal and the Thread ID into Rădulescu2004) are already realized in Rădulescu2004 without the combination of OCP. A POSITA would realize that there is no benefit gained from these combinations.

209. Notwithstanding my opinions regarding the combination of Rădulescu2004 and OCP, the primary indicator of nonobviousness remains true: Rădulescu2004 fails to 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 said at least one communication service identification,” as required by claim 6 of the ‘052 Patent. *Supra* § IX.A.1.ii. OCP fails to disclose the claimed “mapping” step, at least for the reason that OCP fails to disclose a “connection based on a set of connection properties.” Dr. Hassoun does not opine that OCP discloses the “mapping” step. *See* Ex. 1003 at ¶¶ 174-177. Therefore, because neither Rădulescu2004 nor OCP disclose the claimed “mapping ... to a connection based on a set of connection properties,” Rădulescu2004 in view of OCP cannot render obvious claim 6 of the ‘052 Patent.

3. Rădulescu2004 in View of OCP Does Not Render Obvious Claims 1-5 and 7

210. I understand that Dr. Hassoun opines that Rădulescu2004 in view of OCP discloses all of the limitations of claims 1-5 and 7 and that a POSA would have been motivated, with a reasonable expectation of success, to combine Rădulescu2004 with OCP. Ex. 1003 at ¶¶ 178-182. I disagree with Dr. Hassoun's conclusion for claims 1-2, 5, and 7 for the same reasons discussed above for claim 6. *Supra*, §§ IX.A, IX.B.1, IX.B.2.

C. Ground 3: Purported Anticipation by Weber (Claims 1-7)

211. In Ground 3, Dr. Hassoun opines that claims 1-7 of the '052 Patent are anticipated by Weber. Ex. 1003 at ¶¶ 114-157. For the reasons discussed below, I disagree.


212. At the outset, I note that the Board's Institution Decision found that Petitioner did not establish a reasonable likelihood of success in connection with Ground 3. Institution Decision at 43.

1. Weber Does Not Disclose a Single Integrated Circuit (Claims 1-7)

213. Each of the independent claims of the '052 Patent recites an "integrated circuit." Independent claim 1 is directed to an "[i]ntegrated circuit" that comprises a plurality of processing modules (M, S), an interconnect means (N), at least one network interface (NI), and a mapping means (A). Ex. 1001 at 9:25-50.

214. The preamble of independent claim 6 recites “an integrated circuit[] having a plurality of processing modules,” an interconnect means (N), and at least one network interface (NI). Ex. 1001 at 10:7-18. Because the claim specifies that the plurality of processing modules is on the integrated circuit and because the interconnect means (N) couples the plurality of processing modules (M, S), a POSITA would understand that the interconnect means (N) is also on the integrated circuit. The same is true for the network interface (NI).

215. As discussed further below in connection with the “interconnect means,” Dr. Hassoun asserts that the communication pathway between functional block A in sub-system 102 and sub-system 110 (FPGA) that is annotated with red arrows corresponds to the interconnect means:



Ex. 1003 at ¶ 199 (annotating Ex. 1007 at FIG. 1)

216. In FIG. 1 of Weber, the sub-system 102 “is a fully custom integrated circuit.” Ex. 1007, ¶ [0037]. Sub-system 104 “is an application specific integrated circuit (ASIC), which, as is known, is an integrated circuit designated to perform a particular function.” *Id.* The communications bus 112 connects “separate integrated circuit chips.” *Id.* In other words, sub-systems 102, 104, 106, 108, 110 are not on the same integrated circuit or chip in FIG. 1 – they are separate and distinct circuits or chips.

217. Dr. Hassoun opines that Weber suggests that all of the components of system 100 in Figure 1 can be combined on a single integrated circuit. Ex. 1003 at ¶ 185. I disagree. Dr. Hassoun cites the following sentence from Weber: “[s]ystem 100 is an example of a system that may consist of one or more integrated circuits or chips.” Ex. 1007, ¶ [0038]. A POSITA would not understand this to mean that all of the sub-systems 102, 104, 106, 108, 110 can be implemented on a single integrated circuit or chip. Weber also provides that some “embodiments may also include fewer subsystems than the sub-systems shown in system 100.” Ex. 1007, ¶ [0037]. This means that Weber contemplates embodiments with only one of the sub-systems 102, 104, 106, 108, and 110. For example, Weber contemplates an example of the system 110 with only the sub-system 102. As such, a POSITA would understand Weber’s statement that the system “may consist of one or more integrated circuits or chips” to be consistent with the fact that the system 100 may

only include one of the sub-systems (e.g., sub-system 102), in which case the system 100 would only include one integrated circuit or chip. A POSITA would not read this statement in Weber to mean that all of the sub-systems 102, 104, 106, 108, and 110 (or some combination of those sub-systems) can be implemented on single integrated circuit or chip. In other words, a POSITA would understand Weber to mean that system 100 could include one or more integrated circuits, but not that the system 100 (as depicted in Figure 1) can be deployed as a single integrated circuit. Weber never expressly describes or illustrates any embodiment of system 100 where all of the sub-systems are on one integrated chip – they are shown and described as being separate and distinct.

218. Dr. Hassoun also opines that “[b]ecause Weber also suggests that system 100 can be implemented on a single integrated circuit, a POSA would at once envisage that system 100, which includes a plurality of sub-systems 102, 104, 106, 108, and 110, can also be ‘an integrated circuit having a plurality of processing modules (M, S).’ Ex. 1003 at ¶ 191. I disagree at least for the reasons discussed above: Weber does not expressly disclose or suggest that all of the sub-systems 102, 104, 106, 108, and 110 can be implemented on a single integrated circuit. Additionally, there are many technical reasons why a POSITA would not “at once envisage” that Weber discloses a single integrated circuit containing all of the sub-systems 102, 104, 106, 108, and 110 in FIG. 1.

219. A POSITA would also understand that it was, at a minimum, highly impractical and unusual in the relevant timeframe to design and employ a single integrated circuit having all the components of system 100. This is because DRAM 106 is inherently incompatible with the logic processes of the integrated circuits 102 and 104 in FIG. 1 of Weber. There are several reasons for this. First, the CPUs within an integrated circuit are optimized for switching speed and performance using low-latency, high-performance transistor gates, while DRAM is optimized for density and low cost per bit, focusing on optimal design of charged capacitors that store data. Because of DRAM's capacitor-based design, the materials and structure of DRAM is different from the transistors used in CPUs. Integrated DRAM and CPUs on the same chip would compromise either the performance of the CPU or the cost and density of DRAM.

220. Second, DRAM consumes a large amount of space because it stores a significant amount of data. In contrast, CPUs are packed with logic circuits, cores, and cache memory that are designed to maximize processing power within a limited die area. Adding DRAM onto the same chip as an integrated circuit would significantly increase the die size, rendering the chip too expensive and low performing.

221. Third, DRAM consumes less power and generates less heat compared to CPU logic circuits. In contrast, high-performance CPUs generate a lot of heat

and consume substantially more power than DRAM. Thus, a POSITA would understand that integrating DRAM and CPUs on the same chip would introduce thermal management issues, causing both the DRAM's and CPU's reliability and performance to suffer.

222. Fourth, DRAM is typically upgraded more frequently than CPUs because, over time, applications become more complex, thus, demanding more memory space. By decoupling DRAM from the integrated circuit, chip designers can easily increase memory capacity without replacing the CPU. If DRAM were integrated into the same integrated circuit as the CPU, the entire chip would have to be replaced to upgrade memory, thereby reducing flexibility and increasing costs.

223. Indeed, to highlight the difficulties faced in the semiconductor industry of incorporating DRAM on a single integrated circuit, Intel won the Insight Award for Logic in the 2014 TechInsights awards for its Haswell GT3e G82494 integrated circuit which featured an on-chip eDRAM (embedded DRAM). Ex. 2077 at 2. Intel was lauded for implementing the concept of DRAM on an integrated circuit in 2014, more than 10 years after the Weber application was filed.

224. This understanding is supported by Weber. Specifically, Weber's system 100 is designed to "meet the requirements of today's large, complex

electronics systems.” Ex. 1007 at ¶ [0009]. Weber also provides that it is “desirable for a communications system and mechanism to allow sub-systems of a large, complex electronics system to interoperate efficiently.” *Id.* 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.

225. Weber also recognized the need to design communication systems that have a low size or form-factor. Ex. 1007 at ¶¶ [0002]-[0003]. Incorporating all the components of system 100 on a single integrated circuit or chip would result in an integrated circuit or chip that is too large to meet the industry’s requirements in the relevant timeframe.

226. Even today, I am not aware of any commercial integrated circuit that contains each of the sub-systems in FIG. 1 of Weber, i.e., a single integrated circuit that includes a processor, an FPGA, an ASIC, DRAM, and EEPROM.

227. Dr. Hassoun also opines that “[b]ecause Weber also suggests that system 100 can be implemented on a single integrated circuit, a POSA would at once envisage that system 100, which includes a plurality of sub-systems 102, 104, 106, 108, and 110, can also be ‘an integrated circuit having a plurality of processing modules (M, S).’” Ex. 1003 at ¶ 191. I disagree. A POSITA would not “at once envisage” that Weber’s system 100 as shown in Figure 1 can be integrated on

a single integrated circuit because such integrated circuits did not exist in the relevant timeframe. As provided above, a POSITA would understand that integrating every component of Weber's system 100 onto a single integrated circuit would be difficult, and, even if successfully done, the resulting integrated circuit would be lacking in performance and reliability. Even today, where integrated circuit technology has advanced tremendously since the time of Weber's publication, a POSITA would not envisage that all the components of Weber's system 100 would be integrated onto a single integrated circuit. Instead, a POSITA would understand that each respective sub-system (e.g., 102, 104, 106, 108, 110) would be separate integrated circuits connected by the shared-bus 112. This is the only logical conclusion, and it is supported by Weber's disclosure, which provides, "the sub-systems connected by communications bus 112 are separate integrated circuit chips." Ex. 1007 at ¶ [0037].

228. My opinions regarding the lack of a single integrated circuit in Weber would be the same regardless of whether "integrated circuit" is construed as "interconnected circuitry on a chip" (Ex. 1016 at 13) or as "a combination of interconnected circuit elements inseparably associated on or within a continuous substrate" (Ex. 2022 at 59).

229. For these reasons, I disagree with Dr. Hassoun that Weber's system 100, as depicted in Figure 1, can be implemented as a single integrated circuit.

**2. Weber Does Not Disclose the Claimed “Mapping Means”
(Claims 1-5 and 7)**

230. Independent claims 1 and 7 recite a “mapping means.” As discussed above, I understand that Petitioner and Dr. Hassoun have proposed to construe “mapping means” as a means-plus-function term with the corresponding structure being an adapter unit that includes a sequentializing unit and a desequentializing unit. Ex. 1003 at ¶¶ 84-88. This structure is not disclosed by Weber. Dr. Hassoun opines that Weber discloses the claimed structure of the “mapping means.” Ex. 1003 at ¶ 222. I disagree.

231. Dr. Hassoun opines that because Weber discloses that it capable of “multiplexing,” Weber discloses the structure of the claimed “mapping means.” Ex. 1003 at ¶ 225. I disagree. First, Dr. Hassoun fails to opine that Weber contains both a sequentializing unit and a desequentializing unit, as required under Petitioner’s proposed claim construction. Dr. Hassoun opines that Weber’s disclosure of “multiplexing” is analogous to the ‘052 Patent’s disclosure of “sequentializing.” *Id.* But Dr. Hassoun does not specifically identify a structure in Weber that multiplexes the signals sent from Weber’s functional units and that performs the function of the sequentializing unit. *Id.* Additionally, Dr. Hassoun fails to identify any purported desequentializing structure or function in Weber. *Id.*

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.

232. Second, even if Weber did disclose the claimed desequentializing unit (which it does not), a POSITA would understand that Weber does not disclose sequentializing in an integrated circuit. Weber teaches that “multiplex[ing] address, data, and control signals” is “more appropriate for chip-to-chip communications between semiconductor devices.” Ex. 1007 at ¶ [0036]. In contrast, Weber teaches that having “separate command, address, and data wires” is “more appropriate for on-chip communication where wires are relatively less expensive and performance requirements are usually higher.” *Id.* As discussed above, Weber does not disclose a single integrated circuit having all of the sub-systems; but even if it did, Weber teaches that multiplexing is not appropriate for such an embodiment and certainly does not expressly disclose an embodiment with a single integrated circuit that employs multiplexing.

233. Third, Dr. Hassoun opines that the multiplexing would occur “on the shared communication buses 112, 114.” Ex. 1007 at ¶ 225. In other words, Dr. Hassoun opines that multiplexing would occur in the “interconnect means (N),” and not the “mapping means (A).” A POSITA would understand that the “mapping means (A)” in the ‘052 patent is structurally distinct from the functional blocks requesting the communication. *See* Ex. 1001 at FIG. 1, 5:10-13 (“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.”). Because Weber discloses that the structure associated with the multiplexing function is the “shared communications bus 112, 114,” the function of multiplexing alone cannot disclose the “mapping means (A)” under Petitioner’s proposed construction.

3. Weber Does Not Anticipate Independent Claim 6

i. Weber Does Not Disclose the Preamble of Claim 6

234. Dr. Hassoun opines that the “Weber discloses the preamble of claim 6” and addresses the preamble in two pieces. Ex. 1003 at ¶ 185. For the reasons discussed below, I disagree.

a. 6.pre.1 “Method of communication service mapping in an integrated circuit having a plurality of processing modules (M, S)”

235. In discussing limitation [6.pre.1], Dr. Hassoun opines that “the system 100 performs a method of communication service mapping involving thread IDs and CONNIDs” and Dr. Hassoun addresses this limitation in his discussion of limitation [6.d]. Ex. 1003 at ¶ 192. As discussed above, Weber does not disclose an “integrated circuit,” therefore, Weber cannot disclose a “[m]ethod of communication service mapping in an integrated circuit,” for the same reasons.

Supra § IX.C.1.

b. 6.pre.2 “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”

236. Limitation [6.pre.2] recites “at least one communication service identification” that “comprises at least one communication thread or at least one address range.” Dr. Hassoun opines that Weber discloses a communication service identification that comprises at least one communication thread and at least one address range. Ex. 1003 at ¶ 193. I disagree. Weber does not disclose this portion of the preamble.

237. Dr. Hassoun opines that Weber discloses “at least one communication service identification, wherein said at least one communication service identification comprises at least one 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).” Ex. 1003 at ¶¶ 193-198. Specifically, Dr. Hassoun opines that Weber discloses both the thread and the address range. I disagree. Weber fails to disclose a communication

service identification comprising at least one address range for the reasons provided below.

238. Weber does not disclose a communication service identification comprising an address range. Weber discloses a plurality of signals that are carried by a bus. One such signal is the “Addr” signal which “is typically an indication of a particular resource that an initiator functional block wishes to access.” Ex. 1007 at ¶ [0080]. Dr. Hassoun opines that Weber’s “Addr” signal is analogous to the “at least one communication service identification comprises ... at least one address range,” as recited in claim 6 of the ‘052 Patent. Ex. 1003 at ¶ 193. I disagree.

239. Dr. Hassoun’s opinions are based on FIG. 3 of Weber which depicts “a block diagram of an embodiment of a target interface module 900.” Ex. 1003 at ¶ 193 (citing Ex. 1007 at [0046]). Weber’s FIG. 3 relates to the “**target interface module**.” Weber’s FIG. 3 does not disclose the signals that the initiator module (M) sends in its request for a communication service. A POSITA would understand that the communication service request is sent from the master module (M) and that claim 6 requires that the communication service include a communication service identification such as an address range. A POSITA would not understand Weber’s FIG. 3 and its corresponding description to disclose that the master module (M) sends a communication service identification comprising an address

range because Weber's FIG. 3 describes the target interface module, and not the initiator interface module.

240. Dr. Hassoun also opines that "[a] POSA would understand that this [i.e., the decoding of the 'Addr' signal at the target interface] discloses the **capability** of issuing a request based on an address range." Ex. 1003 at ¶ 193 (emphasis added). Dr. Hassoun effectively admits that Weber does not in fact disclose request a communication service including an address range.

241. As explained above, *supra*, § IV.B, it is my understanding that if a prior art reference fails to explicitly disclose a claim element, the prior art reference must inherently disclose the claim element for the prior art reference to anticipate the claim. I also understand that a prior art reference only inherently discloses a claim element if the prior art reference must necessarily include the undisclosed claim element. That Weber has the "**capability**" of issuing a request for a communication service including an address range is insufficient for Weber to inherently disclose issuing a request for a communication service including an address range. I note that Dr. Hassoun makes no attempt to show that Weber must necessarily include the address range.

242. Dr. Hassoun attempts to show that Weber discloses an address range; however, Dr. Hassoun's analysis is confusing at best. Ex. 1003 at ¶ 193. Specifically, Dr. Hassoun opines that:

- Weber discloses that the “Addr” signal is sent by an initiator;
- Weber discloses that the state registers 916 of interface 900 “store, for example, client address base” (Ex. 1007 at ¶ [0046]);
- “A POSA would understand that the client is target module 918”
- “A POSA would understand that an address base indicates that the target has a range of addresses that starts at the address base”
- Weber discloses that “a single target is selected to receive the write data by decoding an external address portion of ADDRE0” (Ex. 1007 at ¶ [0056]);
- “A POSA would understand that ... [the] external address portion of ADDRE0 ... would refer to the [client] address base of the target”;
- “A POSA would understand that ... the remaining portion [of ADDRE0] must indicate internal portions of the address and thus indicating a range of addresses”;
- Therefore, a POSITA would understand that Weber’s write request is based on the alleged address range.

Ex. 1003 ¶ 193 (p. 123-126). There are several problems with Dr. Hassoun’s interpretation of Weber and what a POSITA would glean from Weber. I address each in turn.

243. First, Dr. Hassoun fails to provide any explanation as to why a POSITA would understand that the “client” in the “client address base” refers to the target. *See* Ex. 1003 at ¶ 193. Indeed, the term “client” is used only twice in Weber, each time in the larger phrase “client address base.” *See* Ex. 1007 at ¶¶ [0045], [0046]. A POSITA would understand that if the “client address base” were the “target address base,” Weber would have referred to it as the “target address base.” Weber only refers to the functional units performing the requests (i.e., the “one or more second processing modules (S)” in the ‘052 Patent) as “target(s).” *See generally* Ex. 1007. Moreover, Weber explains that “FIG. 2 is a block diagram on an embodiment of an initiator interface module 800.” *Id.* at ¶ [0043]. The initiator interface module 800 includes “configuration registers 810 [which] store bits that determine the state of module 800, including bandwidth allocation and client address base.” *Id.* at ¶ [0045]. Thus, a POSITA would understand that Weber’s “client address base” “determine[s] the state of [initiator interface] module 800,” and does not refer to the target performing the request.

244. Second, Dr. Hassoun fails to provide any explanation as to why a POSITA would understand that the “address base [in the client address base] indicates that the target has a range of addresses that starts at the address base.” *See* Ex. 1003 at ¶ 193. Indeed, the term “range,” let alone “address range,” is never used in Weber. *See generally* Ex. 1007. Weber does not disclose that the “client

address base” stored in configuration register 410 includes or suggests an “address range.”

245. Third, Dr. Hassoun fails to provide any explanation as to why a POSITA would understand that the “external address portion of ADDRE0 ... would refer to the [client] address base of the target.” Ex. 1003 at ¶ 193. Indeed, as discussed above, the “client address base” “determine[s] the state of [initiator interface] module 800,” and does not refer to the target performing the request. Weber does not disclose that the “external address portion of ADDRE0” refers to the “client address base,” and, in my opinion, this would not be readily apparent to a POSITA reviewing Weber.

246. Fourth, Dr. Hassoun fails to provide any explanation as to why a POSITA would understand that the “the remaining portion [of ADDRE0] must indicate internal portions of the address and thus indicating a range of addresses.” Ex. 1003 at ¶ 193. Weber does not disclose that there must be remaining portions of the ADDRE0 address that are distinct from the external portion of ADDRE0. Similarly, Weber does not disclose that the “remaining portion [of ADDRE0] must indicate internal portions of the address and thus indicating a range of address,” and, in my opinion, this would not be readily apparent to a POSITA reviewing Weber.

247. Fifth, Dr. Hassoun fails to provide any explanation as to why a POSITA would understand that “Weber discloses that its requests can be based on ... an address range.” Ex. 1003 at ¶ 193. Dr. Hassoun is importing into Weber’s disclosure something that is not present or readily apparent to a POSITA.

248. Because Weber does not disclose the claimed “address range,” Weber cannot disclose that the claimed “address range for identifying one or more second processing modules (S).” This is apparent from Weber’s disclosure.

249. Weber describes that an initiator’s request for a read and write transfer identifies the location of the read or write transfer based on the “Addr” signal. Ex. 1007 at ¶ [0052] (“Address lines 228 carry a multi-bit signal driven by an initiator bus owner to specify the address of the object to be read or written during the current transfer.”). A POSITA would understand that “the object to be read or written” includes data to be read from or written to a target. Based on this disclosure, a POSITA would understand that the address signal (“Addr”) does not identify the target itself but identifies an address that may be within the target.

250. Weber also further describes the transfer between initiator and target, which makes clear that the “Addr” signal does not “identify[] one or more second processing modules (S).” *See* Ex. 1007 at ¶ [0052] (“Response lines 232 carry a multi-but signal driven by a target to indicate the status of the current transfer. ... A NULL response indicates that the current transfer is to be aborted, presumably

because **the address does not select any target.**”) (emphasis added). Thus, Weber’s disclosure refers to the address (i.e., the “Addr” signal) “select[ing]” the target, not identifying the target.

251. Weber also describes that the identifier of the target functional block is stored within the target interface module. Ex. 1007 at ¶ [0046] (“[t]arget interface module 900 includes ... state registers in state control block 916 [that] include registers that store, for example, client address base and **an identifier for target functional block 918.**”) (emphasis added). Thus, Weber discloses that the target functional block is identified through the target interface module 900. In other words, the target functional block is not identified by the address signal (“Addr”) sent by the initiator functional block nor the “client address base” stored in the configuration register 410 of the initiator interface module 400. Thus, Weber does not disclose “at least one first of said processing modules (M) requests at least one communication services ... based on ... at least one communication service identification, wherein said at least one communication service identification comprises ... at least one address range, said address range for identifying one or more second processing modules (S)” as recited in claim 6 of the ‘052 Patent.

252. Further, because Weber does not disclose the claimed “address range,” Weber cannot disclose that the claimed “address range for identifying ... a memory region within said one or more second processing modules (S),” as

required by claim 6 of the '052 Patent. Dr. Hassoun does not opine that Weber discloses this limitation. *See* Ex. 1007 at ¶¶ 193-198. Weber does not disclose this limitation.

ii. 6.a “coupling said plurality of processing modules (M, S) by an interconnect means (N) and”

253. Limitation [6.a] recites “coupling said plurality of processing modules (M, S) by an interconnect means (N).” Weber does not disclose this limitation.

254. As explained above, *supra* § VIII.A, the “interconnect means (N)” is a “network” or a “multi-hop interconnect separated by one or more network nodes.” As provided above, a POSITA would understand that there are two critical aspects of a “network” or “interconnect means (N)”: (1) it must implement a “multi-hop” scheme, with “hop” being defined as a router/switch that a packet originating from an arbitrary protocol may traverse and (2) it is an indirect interconnect meaning that the routers and switches within the network need not be directly connected. *Supra*, § VIII.A.

255. I note that Dr. Hassoun relies on two separate and distinct embodiments disclosed in Weber as the alleged “interconnect means (N).” Ex. 1003 at ¶¶ 199-205. Specifically, Dr. Hassoun relies on Weber’s FIG. 1 embodiment and Weber’s FIG. 14 embodiment. *Id.* Neither of these embodiments

disclose an “interconnect means (N)” as claimed by the ‘052 Patent. I will address each embodiment in more detail below.

256. First, Dr. Hassoun relies on Weber’s FIG. 1 embodiment as allegedly disclosing the claimed “interconnect means (N).” Ex. 1003 at ¶ 199. Weber’s FIG. 1 embodiment cannot disclose the claimed “interconnect means” because it is not a “multi-hop interconnect separated by one or more network nodes.” Instead, Weber discloses that “[t]he present invention is a communications system and method for allowing functional blocks or sub-systems of a complex electronics system to communicate with each other through a shared communications resource, such as **a shared bus**.” Ex. 1007 at ¶ [0035] (emphasis added). I understand that the Board, in its Institution Decision, found that “the corresponding structure for the interconnect means should not include a mere bus, switch, and point-to-point wire.” Institution Decision at 18. I agree with this finding. Weber does not disclose the claimed “interconnect means (N)” because Weber makes clear that it merely discloses a “shared bus.” Ex. 1007 at ¶ [0035]. Indeed, the term “network” or “multi-hop” is not used once in Weber. *See generally* Ex. 1007.

257. In particular, Weber’s FIG. 1 discloses a “[s]hared communications bus 112 [that] connects sub-systems 102, 104, 106, 108, and 110.” Weber’s FIG. 1 does not disclose the network for several reasons. First, Weber provides that “the sub-systems connected by communications bus 112 are separate integrated circuit

chips.” Ex. 1007 at ¶ [0037]. Thus, FIG. 1 does not disclose a “method of communication service mapping in **an integrated circuit**, ... comprising the steps of: coupling said plurality of processing modules (M, S) by an interconnect means ...” because the alleged plurality of processing modules (M, S) in Weber’s FIG. 1 are not disposed on the same integrated circuit. There is no disclosure in Weber that the individual components of system 100 can be combined on the same integrated circuit. *Supra* § IX.C.1.

258. Another reason Weber’s FIG. 1 does not disclose the claimed “interconnect means (N)” is because it does not implement the multi-hop scheme. I understand that Dr. Hassoun opines that the multi-hop scheme is implemented in Weber’s FIG. 1. As explained further below, I disagree.




FIG. 1

Ex. 1003 at ¶ 199; Ex. 1007 at FIG. 1 (annotated)

259. As shown in FIG. 1, a plurality of functional blocks A, B, C, D, and E are connected to shared bus 114. The communication from functional block A travels through interface 120, along shared bus 114 until it reaches interface 124. From interface 124 the communication travels along interconnect 115 until it reaches functional block F (which Dr. Hassoun opines is a bridge (Ex. 1003 at ¶ 201)). From functional block F the communication travels along interconnect 112 until it reaches its destination, FPGA 110, specifically, functional block A within FPGA 110. Nowhere in FIG. 1 is there a multi-hop scheme implemented, nor is there a disclosure of a “network node.”

260. Dr. Hassoun opines that “a POSA would understand that F of sub-system 102 is a functional block” and that “sub-system A of block 102 and sub-system A of block 110 are separated by multiple network nodes, including nodes 120, 124, F, and 128.” Ex. 1003 at ¶ 201. I disagree. Elements 120, 124, and 128 in Weber’s FIG. 1 are interfaces. Ex. 1007 at ¶ [0040]. Moreover, Weber discloses that interface 128 is “part of [its] respective functional block or sub-systems,” meaning that it is not a distinct component. *Id.* Dr. Hassoun also opines that functional block F is a bridge, and, thus, a network node. Ex. 1003 at ¶ 201. I disagree. It is irreconcilable to opine that an element is both a “functional block” and a “network node.”

261. A POSITA would understand that the interfaces 120, 124, and 128 or the “bridge” F are not “network nodes.” This is because Weber discloses that the interfaces, the bridge (i.e., a functional block), and the shared bus all communicate in the same protocol. Ex. 1007 at ¶ [0041] (“Functional blocks all communication with one another via shared communications bus 112 or shared communications bus 113 and the protocol of one embodiment.”); *id.* at ¶ [0035] (“a communications protocols allows multiple functional blocks on a single semiconductor device to communicate to each other.”). A POSITA would understand that the interfaces and the “bridge” F communicate in the same protocol as the bus (e.g., the OCP protocol). In contrast, the “interconnect means (N)” is an interconnect in which

data is routed in a network protocol that is abstracted from the initiators and targets. In other words, the “interconnect means (N)” sends data in a **different protocol from that of the initiators and targets**. This is clear from the ‘052 Patent’s specification:

- “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.” Ex. 1001 at 6:30-38;
- “The network interface NI provides a conversion of the packet-based communication of the network N to the higher-level protocol (AHB, DTL, MTL) that the IP modules of FIG. 1 use.” *Id.* at 6:61-64;
- The “mapping is relevant for communication protocols which do not include the notion of connection, but are implemented with an interconnect supporting connections (e.g., the Aethereal network on chip). The advantage is that it enables the use of connections for protocols without the notion of connection, allowing the configuration of differential services for different traffic classes. Backward compatibility of on-chip networks is achieved, by means of the adaptor units/shells to enable existing protocols, such as DTL,

AXI and OCP, to use differential services as offered by a connection based interconnect, such as an on-chip network.” *Id.* at 8:47-58

262. Because each of Weber’s alleged “network nodes” (i.e., the interfaces 120, 124, and 128 and the “bridge” F) run on the same protocol as the initiator and targets modules, a POSITA would understand that Weber does not disclose protocol conversion. Similarly, a POSITA would understand that Weber does not disclose that packets are sent on Weber’s buses. Because there is no conversion from the target/initiator protocol (e.g., OCP) to a network-based protocol (e.g., NTTP), Weber cannot disclose the claimed “interconnect means (N).” Similarly, because Weber does not route packets, Weber fails to disclose a “multi-hop” scheme where a “hop” is defined as a router or switch that a packet originating from an arbitrary protocol may traverse. Weber does not expressly disclose any routers or switches. *See* Ex. 1007.


263. Weber’s FIG. 1 also does not disclose the claimed “interconnect means (N)” because each routing path is directly connected. For example, sub-system A is directly connected to interface 120 via interconnect 115; interface 120 is directly connected to interface 124 via bus 114; interface 124 is directly connected to functional block F via interconnect 115; and functional block F is directly connected to FPGA 110. Ex. 1007 at ¶¶ [0039]-[0041]. In contrast, the “interconnect means (N)” is an indirect interconnect where the functional blocks

communicate via network nodes. Ex. 1001 at 1:50-54 (“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.”). Therefore, a POSITA would understand that Weber’s disclosure of a direct interconnect (e.g., the hierarchal bus) indicates that Weber cannot disclose an “interconnect means (N).”

264. Dr. Hassoun appears to suggest that because Weber has multiple buses (e.g., 112 and 114 in FIG. 1) and certain components between the initiator module and the target module, that means that it is a multi-hop interconnect and therefore distinguishable from the traditional buses discussed in the ‘052 Patent. For all of the reasons discussed above, I disagree. It was not uncommon for traditional bus-based architectures that pre-date the ‘052 Patent to have multiple buses (as opposed to just one). A POSITA would readily understand that those implementations including multiple buses are not a multi-hop interconnect. Similarly, the fact that there are “interfaces” in Weber between the initiator module and the target module does not convert Weber into a multi-hop interconnect.

265. The graphic below shows an example of a multi-bus system that includes three shared buses (left) compared to a NoC (right).

Buses to NoC




Ex. 2073 at 15 (annotated)

The functional blocks (i.e., the boxes in yellow, blue, red, and green) can communicate across and through multiple buses. For example, the green functional block in the lower right corner of the diagram could communication with the yellow functional block in the top left corner of the diagram. No POSITA would call this implementation a multi-hop interconnect. Consistent with this understanding, the presentation containing this diagram expressly contrasts this configuration with an NoC and refers to the NoC as an architectural paradigm shift. Ex. 2073 at 15.

266. The architecture in Weber that Dr. Hassoun points to as an alleged multi-hop interconnect is analogous to the example shown above in that there are multiple buses and a functional block between the target and initiator modules. A POSITA would not consider that configuration to be a multi-hop interconnect.

267. For at least these reasons, Weber's FIG. 1 does not disclose the claimed "interconnect means (N)."

268. Second, Dr. Hassoun relies on Weber's FIG. 14 embodiment as allegedly disclosing the claimed "interconnect means (N)." Ex. 1003 at ¶ 199. I disagree. Weber's FIG. 14 embodiment is not a "multi-hop interconnect separated by one or more network nodes." Dr. Hassoun opines that Weber's FIG. 14 discloses the "multi-hop" scheme because of the annotated figure below:



Ex. 1003 at ¶ 199; Ex. 1007 at FIG. 14 (annotated)

269. In FIG. 14, the communication subsystems are not network nodes, the system performs no hops, and there is no protocol conversion. Weber does not

disclose that communication subsystems 1435, 1440 transmit data in a packetized protocol. Thus, based on Weber's disclosure, a POSITA would understand that there is no protocol conversion in Weber's FIG. 14 and, thus, the communication subsystems 1435, 1440 implement the communication between functional blocks in the same communication protocol as the functional blocks. A POSITA would also understand that the interfaces 1445, 1450, 1455, 1460, 1465, 1470, 1475, 1480 also communicate in the same protocol as the functional blocks. Because the communication subsystems 1435, 1440 do not send packets, Weber's FIG. 14 cannot disclose a network for the same reason as Weber's FIG. 1.

270. I understand that Dr. Hassoun opined for the first time during his deposition that Weber's disclosure of "multiplexing" would indicate to a POSITA that Weber discloses packet-based transmission. Ex. 2068 at 141:19-142:15. I disagree. A POSITA would understand that multiplexing signals has nothing to do with packets.

271. Moreover, Dr. Hassoun opines that the "data flow" disclosed in Weber (and as annotated above) exemplifies the "multi-hop interconnection between modules separated by one or more network nodes." Ex. 1003 at ¶ 203. I disagree. FIG. 17 of Weber describes a step 1715 where "a mapping of the data flow through the channels of the system is first determined." Ex. 1007 at ¶ [0104]. FIG. 18 of Weber describes the "exemplary mapping process." *Id.* at ¶ [0107].




FIG. 18

FIG. 17

Ex. 1007 at FIG. 17 (annotated), FIG. 18 (annotated)

272. With respect to the mapping process, Weber describes that “[f]or each component in a path from initiator to target, e.g., through several interfaces and one or more communication subsystems step 1710, an available channel is chosen ... [and] one could choose the channel that closely matches the performance requirements of the initiator’s data flow.” Ex. 1007 at ¶ [0107]. Based on this disclosure, a POSITA would understand that Weber’s communication subsystems do not choose “an available channel.” Indeed, Weber fails to disclose that the communication subsystems choose the channel for the data flow.

273. For example, Weber describes that “[i]n one embodiment, the system is loaded with a revised mapping and the system is reconfigured accordingly.” Ex. 1007 at ¶ [0111]. Because “the system is loaded with a [] mapping,” the communication subsystems do not choose the channels in the mapping process. Similarly, Weber describes that “[i]n another embodiment, reconfiguration of the system may be accomplished by the system itself.” Ex. 1007 at ¶ [0112]. However, Weber fails to provide any explanation as to which components of the system reconfigure the mapping or choose the channel in the mapping process. Further, Weber describes that “the system may initially be operated using a first mapping that generates an aggregation that does not meet the desired guarantees. However, the system would perform the steps described herein to determine a second mapping of data flows that results in aggregate guarantees that meets the application requirements.” *Id.* at ¶ [0113]. In theme with the previous embodiments, Weber fails to disclose which component of the system “determine[s] a second mapping of flows,” or which component chooses the channels during the mapping process.

274. Because a POSITA would understand that the communication subsystems described in Weber’s FIG. 14 do not choose the available channels of the mapping process, it would be understood that the communication subsystems do not switch or route data. Thus, a POSITA would understand that the

communication subsystems are not network nodes, nor are the communication subsystems the “hop” in the claimed “multi-hop interconnection separated by one or more network nodes.”

275. For at least the reasons described above, Weber does not anticipate “coupling said plurality of process modules (M, S) by an interconnect means (N)” as recited in claim 6 of the ‘052 Patent.

iii. 6.c “controlling the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N) by at least one network interface (NI) associated to said at least one first of said processing modules”

276. Limitation [6.c] recites “controlling the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N) by at least one network interface (NI) associated to said at least one first of said processing modules.” Weber cannot disclose this limitation because Weber does not disclose the claimed “interconnect means (N).” *Supra*, § IX.C.3.ii.

277. Even if Weber discloses the claimed “interconnect means (N),” Weber does not disclose limitation [6.c]. Dr. Hassoun opines that Weber’s interfaces are analogous to the claimed network interfaces, and thus, Weber discloses this limitation. Ex. 1003 at ¶ 209. I disagree.

278. As discussed above, Weber does not perform protocol conversion. *Supra* § IX.C.3.ii. In other words, there are no components disclosed in Weber that

convert the communication protocol of the functional blocks to the communication protocol of the alleged network. The '052 Patent discloses that the network interfaces are responsible for this conversion. *See* Ex. 1001 at 6:30-18 (“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.”); *see also id.* at 6:61-64 (“The network interface NI provides a conversion of the packet-based communication of the network N to the higher-level protocol (AHB, DTL, MTL) that the IP modules of FIG. 1 use.”).

279. A POSITA would understand the term “network interface” refers to a component that interfaces to both a functional block and a network. In order for the network interface to interface to a network, it must communicate in the protocol of that network. In order for the network interface to interface to a functional block, it must communicate in the protocol of the functional block. Therefore, a POSITA would understand that the network interface must translate the protocol of the functional block to the protocol of the network.

280. Dr. Hassoun’s opinion that the interface in Weber “decodes incoming commands” is akin to “provid[ing] a conversion of the packet-based

communication of the network (N) to the higher-level protocol (AHB, DTL, MTL) that the IP modules of FIG. 1 use” is incorrect. Ex. 1003 at ¶ 211 (citing Ex. 1001 at 6:61-64). Decoding incoming commands does not indicate that the communication protocol is converted. Indeed, Weber describes that “Address/command decode block 808 decodes an address on shared communications bus 814 to determine if a write is to be performed to registers associated with the initiator functional block 816.” Ex. 1007 at ¶ [0045]. A POSITA would understand that the command is decoded to determine which target is intended for the communication. In other words, the “decoding” described in Weber is not a protocol conversion; instead, it is merely a way of identifying the target of the command or the type of command (e.g., read and write request).

281. As discussed above, every component of Weber (i.e., the functional blocks, the interfaces, and the buses) communicates in the same communication protocol (e.g., OCP). *Supra*, § IX.C.3.ii. That is not true of the ‘052 Patent. A POSITA would understand that network interfaces of the ‘052 Patent must convert the communication protocol of the functional blocks to the communication protocol of the network in order to “control[] the communication between said at least one first of said plurality of processing modules (M) and said interconnect means (N).”

282. For this additional reason, Weber does not disclose limitation [6.c].

4. Weber Does Not Anticipate Claims 2-5, and 7

283. I understand that Dr. Hassoun opines that Weber anticipates claims 2-5, and 7 for the same reasons discussed with respect to independent claim 6 and independent claim 1. I disagree with Dr. Hassoun's conclusion for claims 2-5 and 7 for the same reasons discussed above for claim 6. *Supra*, §§ IX.C.1, IX.C.3.

D. Ground 4: Purported Obviousness over Weber in view of Rădulescu2004 (Claims 1-7)

284. For Ground 4, Dr. Hassoun opines that claims 1-7 of the '052 Patent are obvious over Weber in view of Rădulescu2004. Ex. 1003 at ¶¶ 232-251. For the reasons provided below, I disagree.

285. I note that Dr. Hassoun does not specifically or separately address dependent claims 2-5 in the section of his declaration addressing alleged obviousness over Weber in view of Rădulescu2004. Ex. 1003 at ¶¶ 232-251.

1. Rădulescu2004 Does Not Cure the Deficiencies in Weber

286. As discussed above, Weber does not disclose or suggest that all of the sub-systems 102, 104, 106, 108, and 110 shown in FIG. 1 can be implemented on a single integrated circuit. *Supra* § IX.C.1. In offering his opinions for Ground 4, Dr. Hassoun does not opine that it would have been obvious to modify Weber (whether in view of Weber alone or in view of Rădulescu2004) so that all of the sub-systems (or even some subset of the sub-systems) would be implemented on a single

integrated circuit, let alone implemented on a single integrated circuit that includes Rădulescu2004's NoC. Ex. 1003 at ¶¶ 232-251.

287. Even if Petitioner or Dr. Hassoun had argued that it would have been obvious to modify Weber so that all of the sub-systems 102, 104, 106, 108, and 110 were implemented on a single integrated circuit (which they did not), I would disagree for the reasons discussed above. *Supra* § IX.C.1.

2. A POSITA Would Not Combine Weber and Rădulescu2004

288. Dr. Hassoun opines that it would be obvious to combine Weber's bus-based interconnect and Rădulescu2004's network-based interconnect so that the resulting combination meets the "interconnect means" requirement in the claims. Ex. 1003 at ¶¶ 232-251. I disagree.

289. Weber makes clear that its objective was to develop a shared-bus interconnect to overcome the limitations of traditional bus-based interconnects. Weber describes "point-to-point interfaces between a sub-system and each peer with which it must communicate" and states that such designs offers "protocol simplicity, guaranteed performance, and isolation from dependencies on unrelated sub-systems" while recognizing that these custom interfaces "are by their nature inflexible." Ex. 1007 at ¶ [0004]. Weber also describes using standardized computer buses, which allow "flexibility in system design." Ex. 1007 at ¶ [0005].

290. Weber recognizes several drawbacks and limitations of traditional buses. For example, Weber states that one “disadvantage of computer buses is that each sub-system or component connected to the bus is constrained to use the protocol of the bus” which then “limits the performance of the sub-system.” Ex. 1007 at ¶ [0006]. Weber also states that “bus systems provide limited capability to preserve order in one transaction stream by supporting ‘split transactions’ in which data from one transaction may be interleaved with data from another transaction in the same stream.” Ex. 1007 at ¶ [0007]. Weber further states that traditional “bus systems do not support true concurrency of operations for a sub-system that can process multiple streams of transactions over a single interconnect.” Ex. 1007 at ¶ [0008].

291. Based on the foregoing examples, Weber recognized that the then “[e]xisting communication approaches thus do not meet the requirements of today’s large, complex electronics system” and that it would be “desirable for a communications system and mechanism to allow sub-systems of a large, complex electronics system to interoperate efficiently regardless of their varying performance characteristics and requirements.” Ex. 1007 at ¶ [0009]. To address these problems, Weber discloses “a communications system and method for allowing multiple functional blocks or sub-systems of a complex electronics

system to communication with each other through a shared communications resource, such as a shared communications bus.” Ex. 1007 at ¶ [0035].

292. Dr. Hassoun identifies several features of Rădulescu2004’s NoC that improves on traditional buses, including scalability, flexibility, fault tolerance, and power efficiency. Ex. 1003 at ¶ 239.

293. Dr. Hassoun asserts that “[a] POSA would have understood that using Rădulescu2004’s NoC [in Weber] would improve the scalability and performance of Weber’s communication system.” Ex. 1003 at ¶ 241. I disagree. Weber already teaches that its shared-bus interconnect “meet[s] the requirements of today’s large, complex electronics subsystems” by allowing the sub-systems “to interoperate efficiently regardless of their varying performance characteristics and requirements.” Ex. 1007 at ¶ [0009]. According to Dr. Hassoun, scalability refers to whether the interconnect is “suitable for systems with a large number of cores or functional blocks.” Ex. 1003 at ¶ 239. Weber expressly says that its system can meet the requirements of large complex systems, i.e., systems with a large number of cores or functional blocks. Ex. 1007 at ¶ [0009]. In other words, Weber teaches that its solution provides the scalability that may have been missing from traditional bus architectures. Dr. Hassoun does not identify a particular reason why a POSITA would believe Weber lacks scalability, and in view of these teachings, a POSITA would understand that Weber already addressed that problem. As such, I

disagree with Dr. Hassoun's conclusion that a POSITA would have been motivated to modify Weber to include Rădulescu2004's NoC to improve scalability.

294. With respect to performance, Weber teaches that traditional buses suffer from performance problems in that they do not allow for concurrent data transmission. Ex. 1007 at ¶ [0006]. Weber goes on to disclose that unlike traditional buses, its solution does provide for concurrent data transmissions. Ex. 1007 at ¶ [0070] ("the present invention supports multiple connections per target functional block"); *id.* at ¶ [0071] ("the present invention supports multiple connections per initiator functional block"). In fact, Weber expressly says that its invention "minimizes communication system area and complexity, while delivering high performance and flexibility." Ex. 1007 at ¶ [0072]. As such, a POSITA would have no reason to modify Weber to include Rădulescu2004's NoC to improve performance through offering concurrent data transmission because Weber already provides this performance benefit.

295. Dr. Hassoun asserts that "a router network like Rădulescu2004's NoC can support more concurrent data transmissions than a bus-based system such as that of Weber." Ex. 1003 at ¶ 241. However, Dr. Hasson does not acknowledge that Weber already provides for concurrent data transmissions and does not explain why Rădulescu2004's NoC would inherently provide more concurrent transmissions than Weber, and I have seen no evidence to this effect. Similarly, Dr.

Hassoun asserts that Rădulescu2004's NoC would "enable a larger amount of data to be communicated between the various modules and blocks in an SoC." Ex. 1003 at ¶ 241. Again, I have seen no evidence that Rădulescu2004's NoC would necessarily enable a larger amount of data to be communicated and Weber expressly states that its system is able to meet the requirements of large complex systems. Ex. 1007 at ¶ [0009]. A POSITA reviewing Weber would not recognize any issue with providing concurrent data transmissions or larges amounts of data and would therefore have no reason to look to Rădulescu2004's NoC for adding those features.

296. Weber also discloses that its shared-bus interconnect provides the flexibility missing from traditional buses. Ex. 1007 at ¶ [0010] ("the present inventions includes a shared communications bus for providing flexible communication capability between electronic sub-systems"); *id.* at ¶ [0072] ("the present invention minimizes communication system area and complexity, while delivering high performance and flexibility"); *id.* at ¶ [0112] ("reconfiguration of the system may be accomplished by the system itself providing even further flexibility and adaptability").

297. Weber also discloses that its shared-bus interconnect is more fault tolerant or reliable compared to traditional buses. A POSITA would understand that Weber's disclosure of BUSY, ReqThreadBusy, RespThreadBusy, and other

signals are intended to improve the shared-bus system's reliability. *See* Ex. 1007 at ¶¶ [0052]-[0054], [0058]-[0068], [0087]-[0089], [0093]. A POSITA reviewing Weber would understand that its solution is fault tolerant and therefore would have no reason to look to Rădulescu2004's NoC to provide fault tolerance. And even if Rădulescu2004's NoC provides better fault tolerance than some traditional buses, I have seen no evidence that it is more fault tolerant than Weber's solution.

298. Weber also discloses that its shared-bus interconnect optimizes power consumption to ensure power efficiency. Ex. 1007 at ¶ [0044] ("These protocols allow for higher transfer efficiencies because independent transactions may use the bus while an initiator waits for a long latency target to return data that has been previously requested by an initiator."); *id.* at ¶ [0054] ("use of a CONNID improve system performance and efficiency."); *id.* at ¶ [0068] ("overall system efficiency is increased, since fewer bus cycles are wasted (as would be the case if target A could not distinguish between connections.)"); *id.* at ¶ [0072] ("The control structures required to support implementation of the present invention ... are simple and require much less area than the data buffering area associated with traditional protocols that do not provide efficient fine interleaving of transfers."). A POSITA reviewing Weber would understand that its solution improves power efficiency relative to traditional buses and therefore would have no reason to look to Rădulescu2004's NoC to improve efficiency. And even if Rădulescu2004's NoC is

more power efficient than some traditional buses, I have seen no evidence that it is more power efficient than Weber's solution.

299. Dr. Hassoun states that "Weber and Rădulescu2004 both have the same general goal of creating an on-chip communication system that will accommodate increasing number of modules on a chip." Ex. 1003 at ¶ 236. As discussed above, Weber discloses that it accomplishes that goal on its own without using an NoC. I disagree with Dr. Hassoun's assertion that "[a] POSA would have been motivated to use Rădulescu2004's NoC ... in place of the shared buses of Weber in order to achieve these known benefits, thereby enabling the same interconnect to be used even as more and more functional blocks are integrated into an SoC" (Ex. 1005 at ¶ 241) because Weber discloses that its shared-bus approach achieves similar known benefits of NoCs that are lacking in traditional bus interconnects. A POSITA would not be motivated to combine Weber with Rădulescu2004 because a POSITA would recognize that both Weber already solves the same known problems in prior interconnects.

300. Dr. Hassoun opines that "NoCs had already become popular prior to the filing of the '052 patent." Ex. 1003 at ¶ 49. The '052 patent recognizes as much. Ex. 1001 at 1:34-36 ("Networks on chip (NoC) have received considerable attention recently as a solution to the interconnect problem in highly-complex chips."). While Weber recognized many of the shortcomings and disadvantages

associated with buses that NoCs can improve upon, Weber nevertheless chose to implement a different bus-based solution rather than a NoC.

301. While Dr. Hassoun’s opinions focus on certain advantages associated with NoCs relative to traditional buses, a POSITA would recognize that there are many disadvantages associated with NoCs relative to traditional buses. For example, the table below lists pros and cons of using a NoC versus a bus. These examples are consistent with a POSITA’s understanding of these issues.

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

Ex. 1021 at 8

302. One of the known disadvantages to using a NoC compared to a bus is that it increases the form factor of the device. Weber recognized that “[e]lectronic computing and communications systems decrease in physical size and cost per function” while increasing in complexity. Ex. 1007 at ¶ [0002]. Similarly, Weber recognized the “challenge” that “modern electronic systems” have “form-factor

requirements.” *Id.* at ¶ [0003]. Thus, a POSITA would understand that Weber’s shared-bus communication system was designed with an effort to reduce the form-factor (i.e., keep the die area size low). Because of this understanding, a POSITA would not be motivated to combine Rădulescu2004’s network into Weber’s SoC because “[t]he network has a significant silicon area,” whereas “[t]he silicon cost of a bus is near zero.” Ex. 1021 at 8. In other words, combining Rădulescu2004’s network with Weber’s SoC would “significant[ly]” increase the form-factor and physical size of the integrated circuit, an issue that Weber recognized needed to be fixed. Therefore, a POSITA would not be motivated to combine Rădulescu2004’s network with Weber’s SoC.

303. A POSITA would understand in the relevant timeframe that NoCs would be more expensive and complicated than Weber’s shared bus. Not only are the plurality of routers in Rădulescu2004’s network more expensive than Weber’s shared-bus, but “[s]ystem designers need reeducation for new concepts” provided by NoCs. Ex. 1021 at 8. This is in contrast to bus-based concepts which “are simple and well understood.” *Id.* A POSITA would realize that it would take less design time and overall expense to create Weber’s shared-bus compared to Rădulescu2004’s network. Because of this, a POSITA would not be motivated to combine Rădulescu2004’s network with Weber’s SoC.

304. One of the reasons why introducing Rădulescu2004's network into Weber would add to the cost and complexity of the system is that a bus-based interconnect like Weber's communicates in a bus-based protocol while a network-based interconnect like Rădulescu2004's communicates in a network-based protocol. This protocol mismatch requires conversion between a bus-based protocol and a network-based protocol. While Rădulescu2004 describes that type of protocol conversion, it would be far more complicated and costly to design such a system compared to using Weber as-is. Because, as discussed above, Weber teaches that it already solves all of the alleged problems identified by Dr. Hassoun, a POSITA would have no reason to introduce additional cost and complexity into a system that already meets the requirements of large complex electronic systems.

3. The Resulting Weber and Rădulescu2004 is Undefined


305. At the outset, Dr. Hassoun fails to specifically identify the resulting combination of Weber and Rădulescu2004. Ex. 1003 at ¶¶ 232-251. Instead, Dr. Hassoun appears to assert that a POSITA would be motivated to make a number of combinations to arrive at the claimed invention. *See id.* For example, Dr. Hassoun opines that one combination would be obvious (i.e., taking Rădulescu2004's network in place of Weber's "bus-based interconnect" (*id.* at ¶ 244)), and also opines that "even if Weber's interface modules were somehow found to not be 'network interfaces' as claimed in [the '052 Patent], it would have been obvious in

view of Rădulescu2004 to employ Rădulescu’s network interface units in Weber, thus meeting this limitation.” *Id.* at ¶ 245. I address each of these permutations below.

i. First Proposed Combination: Weber’s System of Communication Service Mapping with Rădulescu2004’s Network

306. Dr. Hassoun opines that it would have been obvious to a POSITA “to provide and use a system of communication service mapping like that taught by Weber with a NoC such as that taught by Rădulescu2004.” Ex. 1003 at ¶ 235. For this combination, Dr. Hassoun’s opinion appears to be that a POSITA would be motivated to substitute Weber’s bus-based interconnect with Rădulescu2004’s network. Ex. 1003 at ¶ 235. For the reasons provided below, I disagree that a POSITA would have been motivated to do so.

307. As I understand Dr. Hassoun’s opinion, Dr. Hassoun proposes a combination similar to the one represented below.



Ex. 1005 at FIG. 8 (annotated); Ex. 1007 at FIG. 4 (annotated)

308. A POSITA would readily recognize that this proposed combination is untenable because Weber's interfaces 1004 and 1006 communicate in a different communication protocol than Rădulescu2004's network. *Compare* Ex. 1007 at ¶ [0049] ("The bus protocol supports arbitration among many initiator functional blocks and target functional blocks for access to the bus.") *with* Ex. 1005 at 3 § 4 ("The network interface (NI) is the component that provides the conversion of the packet-based communication of the NoC to the higher-level protocol that IP modules use."). In other words, a POSITA would understand that Weber's interfaces 1004 and 1006 receive communications from Weber's functional blocks in a first protocol, and transmit those communications to Weber's bus-based interconnect in the same protocol. This is contrary to Rădulescu2004's network

which transmits the communications in a protocol different than that used by Weber's functional blocks. *See* Ex. 1005 at 3 § 4.

309. A POSITA would not be motivated to swap Weber's bus-based interconnect for Rădulescu2004's network because of the differing communication protocols of Weber's interfaces compared to Rădulescu2004's network. A POSITA would understand that merely inserting a network-based interconnect into a bus-based system would cause the system to be incapable of transmitting communications. The network would not be able to receive the requests issued by the initiator functional blocks and relayed by the interfaces, nor would the interfaces be capable of receiving the request sent over the network in the network communication protocol. Dr. Hassoun does not explain if or how a POSITA would modify Weber's interface modules to perform the protocol conversion that would be necessary to make this proposed combination actually work.

310. As a result, a POSITA would not be motivated to make this combination and would have no expectation of success if it even attempted to make this combination. For these reasons, I disagree that with Dr. Hassoun that a POSITA would consider making such a combination.

311. Nevertheless, even if this combination was made (which it would not be for reasons provided above), Dr. Hassoun opines that "mappings requests to connections according to a communication service identification that includes at

least one address range ... was something a POSA would have been motivated to do in Weber with a reasonable expectation of success, in view of either the teachings of Weber itself (via the incorporated-by-reference U.S. Patent No. 5,948,089).” I disagree.

312. The ‘089 Patent cited in Weber suffers from the same deficiencies as Rădulescu2004. Specifically, the ‘089 Patent provides that “[o]ne example of the use of buffered registers is the assignment of the address map in the system.” Ex. 1008 at 13:22-37. The buffered registers of the ‘089 Patent are located in the “initiator interface module.” *Id.* at Fig. 1. The ‘089 Patent further provides that “[e]ach target interface or module has a locally stored mask and match data value that specifies to which address(es) the target will respond.” *Id.* at 13:22-37. The ‘089 Patent further provides that “[p]roper operation of the system requires that **target address ranges must be uniquely assigned in the system**” and, if addresses are changed, “**new assignments should take place simultaneously in all affected interface modules.**” *Id.* This disclosure makes clear that the “address ranges” referred to in the ‘089 Patent are stored in the initiator interface modules. This is similar to Rădulescu2004’s disclosure of “[t]he address range assigned to a slave is configurable in the narrowcast module.” Ex. 1005 at 4 §_4.2. In other words, the address range is not the communication service identifier; instead, the address range is a variable stored in the interface module associated with the


initiator module. A POSITA would understand that the ‘089 Patent does not disclose a communication service identifier in the form of an address range.

313. I similarly disagree with Dr. Hassoun that this combination, if made, would render claims 1, 6, and 7 obvious. This is because Weber does not teach the “network interface” as disclosed in claims 1-7, *supra* § IX.C.3.iii nor does Weber teach a “mapping means” as disclosed in claims 1 and 7, *supra* § IX.C.2.

ii. Second Proposed Combination: Weber’s SoC with Rădulescu2004’s NoC and Network Interfaces

314. Dr. Hassoun opines that it would have been obvious to a POSITA “to use Rădulescu2004’s NoC (including its network interfaces made up of shells and a kernel) in place of the shared buses of Weber.” Ex. 1003 at ¶ 241. For this combination, Dr. Hassoun’s opinion appears to be that a POSITA would be motivated to swap out Weber’s bus-based interconnect and bus-based interfaces with Rădulescu2004’s network and network interfaces. Ex. 1003 at ¶ 241. For the reasons provided below, I disagree that a POSITA would have been motivated to do so. I also disagree that this combination would render obvious any claim of the ‘052 Patent.

315. As I understand Dr. Hassoun’s opinion, Dr. Hassoun proposes a combination similar to the one represented below.



Ex. 1005 at FIG. 8 (annotated); Ex. 1007 at FIG. 4 (annotated)

Even if a POSITA would make this combination (which they would not, for the reasons discussed above), this combination fails to disclose or render obvious all of the claim limitations.

316. As I discussed above, Rădulescu2004 does not disclose the claimed communication service identification because it does not disclose communication threads or address ranges. *Supra* § IX.A.1.i.b.

317. As discussed above for Ground 3, to make up for the lack of a communication service identification comprising an address range in Rădulescu2004, Dr. Hassoun opines that Weber's initiator functional blocks issue requests that include address ranges as a form of communication service identification. I disagree for the reasons discussed above. *Supra* § IX.C.3.i.b.

Neither Rădulescu2004 nor Weber discloses a communication service identification comprising an address range. *Supra* §§ IX.A.1.i.b, IX.C.3.i.b.

318. To the extent Petitioner and Dr. Hassoun are asserting that it would have been obvious for there to be a communication service identification comprising an address range in the resulting combination of Weber and Rădulescu2004 even though neither reference discloses an address range, I disagree. As discussed in further detail above, the only benefit to using an address range for this purpose that Dr. Hassoun identified was that the address range would identify a target or memory within the target. Ex. 1003 at ¶¶ 249-250. But Rădulescu2004 already discloses that this functionality is possible without there being an address range in the communication service identification. *Supra* §§ IX.A.1.i.b, IX.C.3.i.b, IX.D.3.ii; Ex. 1005 at 4 § 4.2.

319. Further, as discussed above for Ground 3, Dr. Hassoun opines that Weber's initiator functional blocks issue requests that include Thread IDs as a form of communication service identification and that a POSITA would be motivated to incorporate Weber's Thread IDs. Ex. 1003 at ¶¶ 249-250. I disagree. A POSITA would not be motivated to incorporate this feature from Weber into the proposed combination because there would be no way to map the requested communication service to a connection based on a set of connection properties according to the Thread IDs or the address ranges. As discussed above,

Rădulescu2004 does not disclose such a mapping. *Supra* § IX.A.1.ii. As a result, the Thread IDs or address ranges issued as part of the request would be futile.

320. Dr. Hassoun also opines that a POSITA would be motivated to make this combination because a POSITA would recognize that it provides for the “mapping address signals to Rădulescu2004’s connections” which allows “the NoC [to] identify which particular slave module is being targeted by a narrowcast communication service request.” Ex. 1003 at ¶ 250. As explained above, *supra* § IX.A.1.ii., Rădulescu2004 discloses a way to “identify which particular slave module is being targeted by a narrowcast communication service request,” i.e., by the target slave’s address. *See* Ex. 1005 at 4 § 4.2 (“Narrowcast 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.”). Weber also does not disclose the use of address ranges, but discloses a way to identify which particular slave is being targeted by a particular request. Ex. 1007 at ¶ [0056] (“a single target is selected to receive the write data by decoding an external portion of ADDRE0.”). Similarly, the ‘089 Patent, which is incorporated by reference in Weber, does not teach mapping according to an address range. *Supra* § IX.D.3.i. Therefore, in my opinion, this alleged “benefit” of the proposed combination would not motivate a POSITA to make the proposed

combination because this “benefit” is already realized by both Rădulescu2004 and Weber without the use of address ranges.

321. Similarly, Dr. Hassoun opines that a POSITA would be motivated to include Thread IDs in this proposed combination because “by mapping Weber’s threads to Rădulescu2004’s connections, the NoC can perform thread-based ordering and other thread information can be communicated through the NoC” and “would also allow differentiated traffic classes through the connid signals.” Ex. 1003 at ¶ 250. As explained above, *supra* § IX.A.1.ii, Rădulescu2004 discloses ordering and providing differentiated traffic classes without the use of Thread IDs. *See* Ex. 1005 at 4 § 4.2 (“To provide in-order response delivery, the narrowcast must also keep a history of connection identifiers of the transactions including responses ... and the length of these responses. In-order delivery per slave of request messages is already provided by the point-to-point connections.”); *see also id.* at 3 § 4.1 (“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.”). Therefore, in my opinion, this alleged “benefit” of the proposed combination would not motivate a POSITA to make the proposed combination because this “benefit” is already realized by Rădulescu without the use of Thread IDs.

322. As discussed above, Rădulescu2004 does not disclose the “mapping” claimed in the ‘052 Patent because it does not expressly disclose selecting a specific connection based on connection properties assigned to that connection. *Supra* § IX.A.1.ii. Assuming that Rădulescu2004’s network (including the network interfaces containing the adapter units) were substituted into Weber (which I disagree with), then the resulting Weber and Rădulescu2004 combination would suffer from the same deficiency as Rădulescu2004 and would not disclose the claimed mapping. Dr. Hassoun does not assert that it would have been obvious to modify Rădulescu2004 to include this functionality. *See* Ex. 1003 at ¶¶ 232-251. Moreover, Weber does not and cannot disclose the “mapping” because Weber uses a bus-based protocol that does not include the notion of connection in this context. *See* Ex. 1001 at 8:47-51; *see also* Ex. 1007 at ¶¶ [0049] (describing a bus-based protocol that “supports arbitration among many initiator functional blocks and target functional blocks for access to the bus.”), [0044].

323. I similarly disagree with Dr. Hassoun that this combination, if made, would render claims 1, 6, and 7 obvious. This is because Rădulescu2004 does not teach the “mapping” as disclosed in claim 6, *supra* §§ IX.A.1.ii, nor does Weber teach a “mapping means” as disclosed in claims 1 and 7, *supra* § IX.C.2.

X. CONCLUSION

I reserve the right to modify or supplement my opinions, if necessary, based on further review and analysis of the evidence in this case, including review and analysis of information that may be provided to me subsequent to the date of this Declaration.

I declare under penalty of perjury under the laws of the United States that the foregoing is a true and correct reflection of my opinions. This declaration is executed this 30th day of September, 2024.

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

Dr. Nader Bagherzadeh