

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSUNG ELECTRONICS CO., LTD.
Petitioners

v.

NETWORK SYSTEM TECHNOLOGIES, LLC
Patent Owner

Patent No. 7,373,449

Inter Partes Review No. IPR2024-00338

DECLARATION OF DR. ERIK CHMELAR

TABLE OF CONTENTS

I.	INTRODUCTION	4
II.	QUALIFICATIONS AND BACKGROUND.....	5
III.	MATERIALS CONSIDERED	10
IV.	UNDERSTANDING OF LEGAL PRINCIPLES.....	11
A.	Prior Art.....	11
B.	Claim Construction	12
C.	Anticipation.....	13
D.	Obviousness.....	15
V.	PERSON OF ORDINARY SKILL IN THE ART.....	19
VI.	OVERVIEW OF THE TECHNICAL BACKGROUND AND ‘449 PATENT	20
A.	Technical Background	20
B.	Challenges of On-Chip Networks	29
C.	Overview of the ‘449 Patent	36
VII.	THE ‘449 PATENT’S PROSECUTION HISTORY.....	39
VIII.	CLAIM CONSTRUCTION.....	41
A.	The “[first/second] module,” “communication manager” and “resource manager” are separate and distinct (Claims 2, 6, and 9-16).....	42
B.	Samsung’s Litigation Construction for “communication manager” (claims 2, 6, and 9–16).....	46
C.	District Court Constructions.....	48
1.	“Network” (Claims 2, 6, and 9–16)	49
2.	“Integrated Circuit” (Claims 2, 6, and 9–16)	50
3.	“Connection Properties” (Claims 2, 6, and 9–16).....	50
4.	Claim 10’s Preamble	51
5.	“Independently Configurable” (Claim 10).....	51
6.	“Connection” (Claims 2, 6, and 9–16).....	51

7. “At Least one of a Switch and a Router” (Claim 11).....	52
8. “Flow Control” (Claim and 13).....	52
IX. ANALYSIS OF THE ’449 PATENT	52
A. Ground 1: Goossens2002 in View of Drake Does Not Render the Challenged Claims of the ’449 Patent Obvious.....	52
1. Independent Claim 10	52
2. Claims 2, 6, 9, and 11–16.....	77
B. Ground 2: Goossens2003 in view of Drake Does Not Render Claims 1–6 or 9–16 of the ’449 Patent Obvious.....	77
1. Independent Claim 10	77
2. Claims 2, 6, 9, and 11–16.....	101
X. CONCLUSION	101

I. INTRODUCTION

1. I make this declaration based upon my own personal knowledge and, if called upon to testify, would testify competently to the matters stated herein.

2. I have been retained on behalf of Patent Owner Network System Technologies, LLC (“Patent Owner” or “NST”), as an expert consultant to provide this declaration concerning the technical subject matter relevant to U.S. Patent No. 7,373,449 (“the ‘449 patent”) in connection with an *Inter Partes* Review (“IPR”) petition filed by Samsung Electronics Co., Ltd. (“Petitioner” or “Samsung”).

3. I have provided an exemplary, but not exhaustive, summary of materials that I considered or relied upon in preparing this report, besides my personal knowledge and experience, below. The paragraphs below may identify additional materials considered or relied upon not listed in Appendix 2. I reserve the right to rely on additional documents as I become aware of them and to respond to documents relied on by Defendant’s expert.

4. I am being compensated for my time in the amount of \$600 per hour for the time I spend on this matter. I have no financial interest in Patent Owner or the ‘449 Patent, and my compensation is not contingent upon the substance of my declaration, any statements or opinions made, or the outcome of this matter.

5. I understand that at the time of this Declaration, the ‘449 Patent is being asserted by Patent Owner in multiple co-pending district court litigations, including:

- *Network System Technologies, LLC v. Samsung Electronics Co., Ltd. et al.*, Case No. 2:22-cv-00481-JRG (E.D. Tex.) (“the Samsung Litigation”);
- *Network System Technologies, LLC v. Texas Instruments, Inc.*, Case No. 2:22-cv-00482-RWS (E.D. Tex.) (“the TI Litigation”);
- *Network System Technologies, LLC v. Qualcomm Inc. et al.*, Case No. 2:22-cv-01331-DAE (W.D. Tex.) (“the Qualcomm Litigation”);
- *Network System Technologies, LLC v. Lenovo Group Ltd.*, Case No. 2:23-cv-00167-JRG (E.D. Tex.) (“the Lenovo Litigation”);
- *Network System Technologies, LLC v. Ford Motor Co.*, Case No. 2:23-cv-00440-DAE (W.D. Tex.) (“the Ford Litigation”); and
- *Network System Technologies, LLC v. OnePlus Technology (Shenzhen) Co., Ltd.*, 2:23-cv-00520 (E.D. Tex.) (“the OnePlus Litigation”).

II. QUALIFICATIONS AND BACKGROUND

6. In 2000, I received a B.S. degree in Electrical Engineering from Michigan Technological University with the highest distinction (summa cum laude). In 2001, I received an M.S. degree in Electrical Engineering from Stanford University, with a focus on computer networks (EE284 Introduction to Computer

Networks, EE384A Internet Protocols, EE384S Network Architectures, and EE384S Multimedia Networks). In 2004, I received a Ph.D. degree in Electrical Engineering from Stanford University, with a primary focus on very large-scale integrated (VLSI) systems that culminated in a dissertation entitled “The Test and Diagnosis of Field-programmable Gate Arrays, advised by Dr. Edward J. McCluskey, and with conclusion of a secondary focus on computer networks (EE384X Packet Switch Architectures). While studying at Stanford University from 2000–2004, I taught undergraduate and graduate level courses as a Teaching Assistant (EE284 Introduction to Computer Networks, EE384A Internet Protocols and Standards, Digital Design Laboratory EE121, and E40 Introduction to Electronics), and I was employed as a Research Assistant in the Center for Reliable Computing. In 2011, I received a second B.S. degree in Chemical Physics from Saginaw Valley State University in 2011 with the highest distinction (summa cum laude).

7. During my undergraduate and graduate studies, I was additionally employed at several product-technology companies. From 1999–2002, I was employed as a Manufacturing Engineer at Cisco Systems, Inc, where I oversaw functional testing of layer-2 and layer-3 network switches and determined root-cause failures thereof. During the summers of 2002 and 2003, I was employed as a

Product Test Engineer intern at Xilinx, Inc., where I researched and developed software tools to detect and diagnose physical hardware defects in FPGAs.

8. After earning my Ph.D., I was employed at LSI Corp. (now Broadcom, Inc.) from 2004–2012 as a Staff Software Engineer in the Advanced Technology Development department, where I conducted R&D for very large-scale integrated (VLSI) systems and serializer-deserializer (SerDes) circuits for high-speed data communications, for which I was granted approximately 20 U.S. patents. During this period of employment at LSI Corp., I also taught at several Universities. From 2006–2008, I was a consulting Assistant Professor in the Electrical Engineering and Computer Science department at Stanford University, where I advised PhD students in the Center for Reliable Computing on the test and reliability of VLSI systems, and created and taught a PhD-level course on clock and data recovery (CDR) for digital communications (EE391 Clock and Data Recovery Architectures). From 2009–2012, I was an Adjunct Professor in the Electrical and Computer Engineering department at Saginaw Valley State University, where I taught several undergraduate courses in data communications and computer hardware (ECE455 Data Communications, ECE368 Computer Engineering, ECE368L Computer Engineering Lab, ECE216 Digital Circuits, and ECE216L Digital Circuits Lab).

9. From 2012–2015, I was employed at Stryker Instruments as a Senior Principal Engineer in the Advanced Development department, where I conducted and supervised R&D of computationally intelligent medical instruments and surgical technologies, and I delivered working prototypes for bio-photonic devices and boundary-constraint surgical robotics.

10. From 2015–2018, I was employed at Interdigital, Inc. as a Senior Manager of Innovation Partners, where I conducted and supervised R&D of 5th generation wireless (5G) network-enabled technologies including the Internet of Things, connected/autonomous vehicles, and augmented/virtual reality infrastructure, and I managed scientists and engineers to create over 100 new-invention disclosures for patent protection.

11. From 2019–present, I have provided consulting services to a major patent-licensing corporation as a Technical Expert in the areas of semiconductor circuits, microchips, and data communications. My consulting services have included determining whether and how specific patent claims within the aforementioned areas read on various products or product features. I analyze patents, technical specifications, product white papers, and product marketing materials to produce detailed claim charts used by attorneys to formulate infringement opinions and licensing decisions. I have also performed technology

market analysis and developed detailed technology landscapes and roadmaps in the aforementioned areas.

12. I have written technical publications dating back to 2003 on topics in the areas of semiconductor circuits, microchips, and data communications. I have written technical articles published in peer reviewed professional journals, trade journals, and proceedings of peer-reviewed conferences. I have presented at several technical symposia and technical conferences on topics in the aforementioned areas.

13. I am or have been a member or senior member of several professional organizations, including the Institute of Electrical & Electronic Engineers (IEEE), the Association for Computing Machinery (ACM), the American Society of Mechanical Engineers (ASME), the International Society of Optics & Photonics (SPIE), and the American Chemical Society (ACS).

14. I served as a Program Evaluator for the Accreditation Board of Engineering and Technology (ABET) from 2013–2018. As a Program Evaluator, I was responsible for determining ABET accreditation for university engineering programs categorized as Electrical Engineering programs and Computer Engineering programs.

15. I served as an Industry Mentor for National Science Foundation (NSF) I-Corps program in 2013, and served as an Instructor for the same in 2015. In these

roles, I provided industry and technical guidance to engineers, inventors, and entrepreneurs launching technology startups.

16. I hold approximately 30 U.S. Patents, many of which concern technologies in the areas of semiconductor circuits, microchips, and data communications.

17. Additional details on my background and experience can be found in my *curriculum vitae*, which I understand is being submitted as Ex. 2002.

III. MATERIALS CONSIDERED

18. In forming the opinions set forth in herein, I have considered and relied upon my education, knowledge of the relevant field, and my experience. I have also reviewed and considered the ‘449 Patent (Ex. 1001), the ‘449 Patent’s file history (Ex. 1002), and at least the following additional materials:

- Samsung Electronics Co., Ltd.’s Petition for *Inter Partes* Review of the ‘449 Patent (Paper 2);
- Declaration of Dr. Todd Mowry (Ex. 1003);
- K. Goossens, et al., *Networks on Silicon: Combining Best-Effort and Guaranteed Services* (Ex. 1005, “Goossens2002”);
- U.S. Patent No. 5,461,661 (Ex. 1006, “Drake”);
- K. Goossens, et al., *Guaranteeing the Quality of Services in Networks on Chip* (Ex. 1007, “Goossens2003”); and

- Any other documents or references cited or referred to in this Declaration.

IV. UNDERSTANDING OF LEGAL PRINCIPLES

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

A. Prior Art

20. I have a general understanding of prior art and priority dates based on my experience with patents, my career experience, and my discussions with counsel.

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

22. 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 Petitioner has the burden of proving that any particular reference or product usage or offer for sale is prior art.

B. Claim Construction

23. I am informed and understand that claim construction is a matter of law.

24. I am informed and understand that the specification, which includes the claims, is the single best guide as to the meaning of claim terms.

25. I am informed and understand that the relevant inquiry in claim construction is the question of how a POSITA would have understood the claim terms at the time of the invention, in light of the patent specification and prosecution history, and a POSITA is deemed to have read the claim terms in the context of the entire patent and its prosecution history.

26. I understand that the patent specification may reveal a special definition given to a claim term by the patentee that differs from the plain and ordinary meaning it would otherwise have to a POSITA. In such cases, I understand that the patentee's definition usually controls.

27. I understand that the prosecution history of a patent can inform the meaning of some claim language and that the prosecution history must be considered when construing the claims.

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

29. I understand that one of ordinary skill in the art is presumed to be a person with at least a particular level of skill and knowledge in a certain field or industry, who is capable of understanding and practicing the technology described in the patent at issue. This person is assumed to have ordinary, not extraordinary, skill.

C. Anticipation

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

31. I understand that the anticipation analysis is a two-step process. The first step is to determine the meaning and scope of the asserted 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.

32. 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.

33. 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.

D. Obviousness

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

35. 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.

36. 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.

37. 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.

38. 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.

39. 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.

40. 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.

41. 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.

42. 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.

43. 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.

44. 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.

45. 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.

V. PERSON OF ORDINARY SKILL IN THE ART

46. 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.

47. I understand that in the TI Litigation, the Court adopted the following identification of 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 as at least two years of work experience in the field of SoC and NoC technologies.” I met and exceeded this level of skill as of around 2002 and have applied this definition of a POSITA throughout this Declaration.

48. I understand that Petitioner and Dr. Mowry assert that a POSITA “would have had at least a Bachelor’s degree in computer science, electrical engineering, computer engineering, or an equivalent technical degree, and two or more years of experience with system-on-chip, interconnect, or networking technologies.” Petition, 14; Ex. 1003, ¶¶ 59-62. I met and exceeded this level of

skill as of around 2002. While I have applied the definition of a POSITA adopted by the court in the TI Litigation, my opinions would not change if the Board were to adopt Petitioner and Dr. Mowry's proposal.

VI. OVERVIEW OF THE TECHNICAL BACKGROUND AND '449 PATENT

A. Technical Background

49. A network-on-chip ("NoC") is a data communication network on a microchip. In some ways, NoCs are like traditional computer networks, such as local area networks ("LANs") and/or wide area networks ("WANs"). For example, both traditional computer networks and NoCs interconnect devices or components: traditional computer networks commonly interconnect discrete devices, such as computers, printers, servers, and so on, whereas NoCs commonly interconnect integrated components, such as processors, cores, caches, memories, input/output (I/O) interfaces, peripherals, and so on. However, implementation of a NoC on a SoC presents many unique challenges, as discussed in more detail below. *Infra*, § VI.B. The components interconnected by a NoC are typically integrated on a very large-scale integrated ("VLSI") circuit, such as a system-on-chip ("SoC").

50. An SoC is an integrated circuit in which most or all components of a system are integrated, including memories, processors, high-speed I/O interfaces, and dedicated application-specific logic. *See, e.g.*, EX. 2003 at 1; *see also* EX. 2004 at 4–5. By the early 2000s, certain limitations of statically wired interconnections

between components of an SoC were known, and as a result, NoCs were already emerging. *See, e.g.*, Ex. 2004 at 7–8, *see also* Ex. 2005. Rather than considering on-chip interconnections as point-to-point wires, they can be abstracted as communication channels over which to transmit data within a “quality-of-service” setting. *See id.*

51. Before describing details of a NoC, it is important to describe how components can be interconnected in an SoC that does not implement a NoC. Interconnections, sometimes called interconnects, traces, wires, and signal lines, are electrical connections between a first component and a second component. Interconnects that carry similar signals are grouped together into buses, for example interconnects that carry address signals are grouped into an address bus and interconnects that carry data signals are grouped into a data bus. *See* Ex. 2006 at 4. A bus can be considered a direct interconnection between components (e.g., the path between a first component and a second component is always the same), where the interconnection may be dedicated or shared. A dedicated bus comprises point-to-point interconnection between components, typically between two components. A shared bus comprises interconnections that can be shared between many components, and requires some type of bus arbitration to resolve conflicts when multiple components attempt to access the shared bus simultaneously.

52. There are several types of bus arbitration methods for shared buses, including centralized and distributed (sometimes called decentralized). In centralized arbitration, a single device, known as a bus controller or bus arbiter, is responsible for managing access to the bus. In distributed arbitration, components compete for access to the bus by sending a request signal to the bus arbiter and waiting for a grant signal. In some distributed arbitration methods, e.g., decentralized, each component has its own priority level, and the component with the highest priority is given access to the bus. A component capable of originating a bus transaction (e.g., requesting control of the bus) is a master component, and a component that cannot originate a bus transaction (e.g., only listens on the bus for communications intended for it) is a slave component. In large and/or complex SoCs, such shared buses interconnecting master devices and slave devices via bus arbitration may not be sufficiently scalable nor provide enough bandwidth. *See Ex. 2007 at 7, 235.*

53. In particular, with increasing quantities of components integrated in SoCs, increasing transistor densities, and increasing clock frequencies, it becomes very difficult to operate all the components synchronously within a same clock domain. This creates a need for a modular approach, where the processing system of the SoC comprises a plurality of relatively independent complex modules. In conventional processing systems, the system modules usually communicate with

each other via a shared bus as described above. But as the number of modules increases, shared-bus communication is no longer practical at least because: (1) the large quantity of modules creates too high of an electrical load on the shared bus, and (2) the shared bus forms a communication bottleneck because only one component can control the bus at a time. Consequently, some SoCs implement shared buses that utilize crossbar switches, also called crossbars, cross-point switch, matrix switch, or switch matrix, to achieve higher throughput, with STBus, IBM CoreConnect, and ARM AMBA being several examples of bus standards that may utilize crossbars. *See* Ex. 2007 at 320. A crossbar has a plurality of input lines and output lines that form a crossed pattern of intersecting lines between which a connection may be established by closing a switch located at a given intersection. An n -input, m -output crossbar, i.e., an $n \times m$ crossbar, is the equivalent to a set of m multiplexers (one per output port), each having n inputs (one per input port). *See* Ex. 2007 at 332–333. However, a shared bus that utilizes crossbars still has disadvantages. For example, the physical implementation of complex SoCs require large and complex crossbars that are difficult to place-and-route during physical design on the SoC in the areas between the components that implement functional logic of the SoC. *See* Ex. 2008 at 3. An on-chip communication network, e.g., a NoC, is an effective way to overcome some of the aforementioned disadvantages.

54. A NoC interconnect paradigm radically changes the on-chip communication compared to direct interconnects, such as buses or switches. This is because of the multi-hop nature of a network, where components are not directly connected, but separated by one or more network nodes. This is in contrast with buses, where components are directly connected as described above. Several implications of this change include arbitration methods (which can no longer be centralized), and communication properties (e.g., ordering, flow control, and so on). In an SoC that implements a NoC, certain design-specific wiring structures of the SoC, typically the longer “global” wires that interconnect components, are replaced with a general-purpose on-chip communication network. Ex. 2005 at 2. Instead of connecting these top-level modules by routing dedicated wires, they are connected to a network that routes packets between them. *See id.* In doing so, these network-connected modules, or components, can be thought of as “network clients” that communicate with each other through communication channels of a network that provide one or more levels of service, e.g., as may be achieved by various communication properties. *See id.*

55. The 2001 paper by Dally and Towles provides a basic example of a NoC implementation. *See* Ex. 2005, at 2–3. It describes dividing a chip into 16 “tiles,” where each tile comprises client logic, such as processors, DSPs, peripheral controllers, memory subsystems, and so on. The client logic within a tile (i.e.,

nearby logic) may communicate via direct interconnections, but client logic in different tiles (i.e., far-apart logic) must communicate only over the network. In other words, there are no top-level (global) bus-based interconnections, only interconnections provided by the network. Figure 1 of the Dally and Towles paper is reproduced below. *See* Ex. 2005 at 2.




Figure 1 Partitioning the die into module tiles and network logic

56. Each tile interfaces with the network over an input port that is used to insert packets into the network and an output port over which the network delivers packets to the tile (where input and output here refer to packets going in and out of the network, not the tile). Each packet includes a virtual channel field that specifies which virtual channel it may be routed on, to provide different levels of service, and route field to specify a route through the network, i.e., the hops taken from source tile to destination tile. Such a port interface can present a simple and reliable level of service to the network clients. Ex. 2005 at 3.

57. Logic local to the network clients can layer higher-level services on top of this interface. For example, this local logic could present a memory read/write service, a flow-controlled data stream, or a logical wire to the client. Local logic can also provide a translation from a destination node to a route. *See* Ex. 2005 at 3.

58. The router needed at each tile consists of five input controllers (one for each direction and one for input from the tile) and five output controllers. The router for each tile is distributed across the four edges of a tile with the input and output controller for a given direction (e.g., west) located on that edge of the tile. Figure 2 of the Dally and Towles paper, reproduced below, shows the west input controller and its connection to four output controllers (one for each of the other directions and one for the tile). *See* Ex. 2005 at 3.




Figure 2 West input connections to output controllers

59. Details of an input controller and output controller are shown in Figure 3 of the Dally and Towles paper, reproduced below. Each input controller

has an input buffer and input state logic for each virtual channel. When a head flit (flow control digit, e.g., part of a packet) arrives at each tile, the input controller strips the next entry off the route field and uses these two bits to select one of four output ports. The flit then arbitrates with the other virtual channels on that input port and, if it wins the arbitration, is forwarded to the output port. This arbitration and forwarding takes place in parallel with allocating a virtual channel and checking available buffer space to reduce latency. The output port provides a single stage of buffering for each input port connection. The flits in these buffers arbitrate for the link to the input controller on the next tile. Credits for buffer allocation are piggybacked on flits travelling in the reverse direction. *See* Ex. 2005 at 4.




Figure 3 Input (top) and output controllers

60. As indicated earlier, the communication between modules (e.g., components, network clients, etc.) are performed over connections. Some examples of connection properties, which are sometimes called communication properties, are provided below.

The connection properties may include 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, or data delivery.

Ex. 1001 at 8:62-9:3.

61. With such connection properties, on-chip networks are typically able to provide varying levels of communication service, which may be referred to as a quality of service (QoS) levels. “The QoS is a very important feature in the interconnect infrastructures because it provides a regulation mechanism allowing specification of guarantees on some of the parameters related to the traffic.” Ex. 2007 at 325. Traffic here refers to a flow of data between a first module and a second module, for example, a series of data packets transmitted from the first module to the second module. With QoS, different traffic classes can be supported, such as:

- Real time and low latency (RTLL)—Traffic flows that require the lowest possible latency.

- Guaranteed throughput (GT)—Traffic flows that must maintain their throughput over a relatively long time interval.
- Guaranteed bandwidth (GBW)—Traffic flows that require a guaranteed amount of bandwidth over a relatively long time interval.
- Best effort (BE)—Traffic flows that do not require guaranteed latency or throughput but have an expectation of fairness.

Ex. 2007 at 326.

B. Challenges of On-Chip Networks

62. The issues and problems of implementing a multi-hop on-chip network interconnect are complex. Prior to the invention of a network-on-chip interconnect, functional blocks of an integrated circuit were typically interconnected using buses or crossbars. However, both buses and crossbars have respective drawbacks, as discussed in more detail below.

63. Buses, such as ARM’s Advanced Microcontroller Bus Architecture (AMBA), were an early solution for interconnecting on-chip functional blocks, e.g., initiators and targets. However, “the integration of multiple cores within chips” caused “too many initiators trying to access different targets simultaneously” which “created bottlenecks on the bus.” Ex. 2009 at 3. Buses also had “latency issues.” *Id.* As a result, “[t]he industry wanted to enable concurrent access and more overall system data throughput.” *Id.*

64. Following the use of buses, crossbars were introduced as a next interconnect solution. While crossbars cure some of the deficiencies caused by buses, “crossbars require complex control logic to keep track of outstanding transactions” and require “wide data paths for the initiators and targets that need the most throughput as well as signaling for address and control information.” Ex. 2009 at 3. “To support continued system scaling, crossbars were cascaded with bridges between fabrics” which “carried a significant cost in silicon area” and “limited clock frequencies and added significant latency to transactions.” *Id.* at 4. Moreover, “during the 2000s, the initiators and targets became so numerous and widely distributed in the physical floorplan of a chip that the crossbar became a physical wiring nightmare.” *Id.*

65. Due to the drawbacks of buses and crossbars, “Network-on-Chip (NoC) technology emerged as a solution to the wiring problem.” Ex. 2009 at 5. NoC technology packetizes address and data information on the same wires, thereby allowing transactions to be sent over a smaller number of wires compared to prior interconnect solutions “while maintaining high quality of service for each transmission.” *Id.* NoC technology also contributed to decreased die density by “distributing the interconnect logic throughout the chip rather than having bridges as chokepoints.” *Id.*

66. The ‘449 Patent describes the distinction between prior interconnects (e.g., busses and crossbars) and a multi-hop interconnect.

NoCs help resolve the electrical problems in new deep-submicron technologies, as they structure and manage global wires. At the same time they share wires, lowering their number and increasing their utilization. ... NoCs also decouple computation from communication, which is essential in managing the design of billion-transistor chips. NoCs achieve this decoupling because they are traditionally designed using protocol stacks, which provide well-defined interfaces separating communication service usage from service implementation.

Ex. 1001 at 1:38–49. Moreover, using multi-hop interconnects “radically changes the communication when compared to direct interconnects, such as buses or switches.” *Id.* at 3:46–57 (“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. The implications of this change reside in the arbitration (which must change from centralized to distributed), and in the communication properties (e.g., ordering, or flow control.)”).

67. Specific examples of some differences between a multi-hop interconnect and a bus include: (1) media arbitration; (2) data routing; (3) data formatting; (4) latency; and (5) flow control for input buffers. Regarding media arbitration, “[i]n a bus, master modules request access to the interconnect, and the arbiter grants the access for the whole interconnect at once,” whereas in a NoC,

“the arbitration is distributed, because it is performed in every router, and is based only on local information.” Ex. 1001 at 5:25–42.

68. Regarding data routing, “[f]or a bus, the command, address, and data are broadcasted on the interconnect. They arrive at every destination, of which one activates based on the broadcasted address, and executes the requested command,” whereas in a NoC, “it is not feasible to broadcast information to all destinations, because it must be copied to all routers and network interfaces. This floods the network with data. The address is better decoded at the source to find a route to the destination module.” Ex. 1001 at 5:43–54.

69. Regarding data formatting, “[i]n modern bus interfaces the data format is defined by separate wire groups for the transaction type, address, write data, read data, and return acknowledgements/errors (e.g., VCI, OCP, AMBA, or CoreConnect),” whereas in a NoC, “there is typically no distinction between different kinds of data. Data is treated uniformly, and passed from one router to another.” Ex. 1001 at 6:8–25.

70. Regarding latency:

[L]atency is caused by two factors: (a) the access time to the bus, which is the time until the bus is granted, and (b) the latency introduced by the interconnect to transfer the data. For a bus, where the arbitration is centralized the access time is proportional to the number of masters connected to the bus. The transfer latency itself typically is constant and relatively fast, because the modules are linked directly. However, the speed of transfer is limited by the bus speed, which is relatively low. In a NoC, arbitration is performed at each router for the following link.

The access time per router is small. Both end-to-end access time and transport time increase proportionally to the number of hops between master and slave. However, network links are unidirectional and point to point, and hence can run at higher frequencies than buses, thus lowering the latency. From a latency prospective, using a bus or a network is a trade off between the number of modules connected to the interconnect (which affects access time), the speed of the interconnect, and the network topology.

Ex. 1001 at 5:55–6:7.

71. Multi-hop interconnects also differ from buses with respect to flow control for input buffers.

For buses, the source and destination are directly linked, and, destination can therefore signal directly to a source that it cannot accept data. This information can even be available to the arbiter, such that the bus is not granted to a transaction trying to write to a full buffer. In a NoC, however, the destination of a transaction cannot signal directly to a source that its input buffer is full. Consequently, transactions to a destination can be started, possibly from multiple sources, after the destinations input buffer has filled up. If an input buffer is full, additional incoming transitions are not accepted, and stored in the network. However, this approach can easily lead [sic] to network congestion, as the data could be eventually stored all the way to the sources, blocking the links in between.

Ex. 1001 at 6:48–62.

72. The ‘449 Patent also emphasizes that “NoC’s premises are different from off-chip networks, and, therefore, most of the network design choices must be reevaluated” because “[o]n-chip networks have different properties (e.g., tighter link synchronization) and constraints (e.g., higher memory cost) leading to different design choices, which ultimately affect the network services.” Ex. 1001 at

1:66–2:5. Another distinction between an on-chip network compared to off-chip networks is that “[a]n off-chip network interface usually contains a dedicated processor to implement the protocol stack up to network layer or even higher, to relieve the host processor from the communication processing;” however, “[i]ncluding a dedicated processor in a network interface is not feasible on chip, as the size of the network interface will become comparable to or larger than the IP to be connected to the network.” Ex. 1001 at 2:16–25.

73. Further, data sent via a NoC may be subject to out-of-order arrival, diverse routing paths, and retransmission after dropping.

In a network, data sent from a source to a destination may arrive out of order due to reordering in networks nodes, following different routes, or retransmission after dropping. For off-chip networks out-of-order data deliver is typical. However, for NoCs where no data is dropped, data can be forced to follow the same path between a source and a destination (deterministic routing) with no reordering. This in-order data transportation requires less buffer-space, and reordering modules are no longer necessary.

Ex. 1001 at 3:8–17.

74. On-chip networks also differ from off-chip networks with respect to time-related guarantees.

Off-chip networks typically use packet switching and offer best-effort services. Contention can occur at each network node, making latency guarantees very hard to offer. Throughput guarantees can still be offered using schemes such as rate-based switching or deadline-based packet switching, but with high buffering costs. An alternative to provide such time-related guarantees is to use time-division multiple access (TDMA) circuits, where every circuit is dedicated to a network

connection. Circuits provide guarantees at a relatively low memory and computation cost. Network resource utilization is increased when the network architecture allows any left-over guaranteed band width to be used by best-effort communication.

Ex. 1001 at 3:34–46.

75. On-chip networks can also provide more reliable communication compared to off-chip networks.

A consequence of the tight on-chip resource constraints is that the network components (i.e., routers and network interfaces) must be fairly simple to minimize computation and memory requirements. Luckily, on-chip wires provide a reliable communication medium, which can help to avoid the considerable overhead incurred by off-chip networks for providing reliable communication. Data integrity can be provided at low cost at the data link layer. However, data loss also depends on the network architecture, as in most computer networks data is simply dropped if congestion occurs in the network.

Ex. 1001 at 2:43–53.

76. On-chip networks also differ from off-chip networks because of their constraints and synchronization.

Typically, resource constraints are tighter on chip than off chip. Storage (i.e., memory) and computation resources are relatively more expensive, whereas the number of point-to-point links is larger on chip than off chip. Storage is expensive, because general-purpose on-chip memory, such as RAMS, occupy a large area. Having the memory distributed in the network components in relatively small sizes is even worse, as the overhead area in the memory then becomes dominant.

Ex. 1001 at 2:6–15.

C. Overview of the '449 Patent

77. The '449 Patent relates to “an integrated circuit having a plurality of processing modules and a network arranged for providing connections between processing modules and a method for exchanging messages in such an integrated circuit.” Ex. 1001 at 1:7–11.

78. For example, FIG. 3 of the '449 Patent shows an embodiment that comprises “a communication manager CM and a resource manager RM.” Ex. 1001 at 9:66–10:3. The “communication manager CM and the resource manager RM are arranged in the router network RN, but they can also be arranged in one or some of the network interfaces NI.” *Id.* at 10:3–6.

79. “The communication manager request [sic] a connection with a set of properties between two modules from the resource manager RM” and the “resource manager RM allocates the necessary resources and enquires if such a connection based on these resources is possible.” Ex. 1001 at 10:19–26. “After enquiring of the available connection properties, the resource manager RM reserves the available connection, i.e. the required resources, and responds to the communication manager CM which connection is available, i.e. which connection properties or the required resources thereof are available for the desired channels.” *Id.* at 10:41–46. If the communication manager accepts the offered connection, “the resource manager RM sets the connection properties, and establishes a

connection with accepted properties between said initiator and target modules.” *Id.* at 10:51–57. After the initiator and the target modules complete the transaction, “the communication manager CM issues a request to said resource manager RM to reset the connection or the connection properties.” *Id.* at 10:57–61.




FIG.3

Ex. 1001 at FIG. 3

80. “The connections ... must be first created or established with the desired properties before being used,” which “may result in resource reservations inside the network (e.g., buffer space, or percentage of the link usage per time unit.” Ex. 1001 at 10:62–66. “To allow more flexibility in configuring connections, and, hence, better resource allocation per connection, the outgoing and return parts of connections can be configured independently.” *Id.* at 11:4–7. Independently configuring connections allows for “a different amount of buffer space [to] be allocated in the NIPs at the master and salves, or different bandwidths can be reserved for requests and responses.” *Id.* at 11:7–10. As a result, “the connection

properties of request part of a connection can be best effort, while the connection properties of response can have guaranteed throughput (or vice versa).” *Id.* at 11:22–26.

81. Thus, ’449 Patent discloses an integrated circuit including an on-chip network where resource managers of the chip are configured to provide a communication channel between modules that have connection properties, where connection properties between different communication channels are independently configurable.

82. I understand that Petitioner is challenging claims 1–6 and 10–16 of the ’449 Patent in the Petition.

83. The ’449 Patent recites two independent claims: claim 1 and claim 10. The exact language of asserted independent claim 10 is reproduced below:

10. [preamble] Method for exchanging messages in an integrated circuit comprising a plurality of modules, the messages between the modules being exchanged over connections via a network, wherein said connections comprises a set of communication channels each having a set of connection properties, any communication channel being independently configurable, wherein said connection through the network supports transactions comprising at least one of outgoing messages from the first module to the second module and return messages from the second module to the first module and further comprising the steps of:

[a] the first module issuing a request for a connection with the second module to a communication manager, wherein the request comprises desired connection properties associated with the sets of communication channels;

[b] the communication manager forwarding the request to a resource manager;

[c] the resource manager determining whether a target connection with the desired connection properties is available;

[d] the resource manager responding with the availability of the target connection to the communication manager; and

[e] the target connection between the first and second module being established based on the available properties of said communication channels of said connection.

84. Challenged claims 11, 12, 13, 14, 15, and 16 all depend from independent claim 10. To the extent that the dependent claims are relevant to my opinions, the specific dependent claims will be reproduced and analyzed accordingly.

VII. THE ‘449 PATENT’S PROSECUTION HISTORY

85. The ‘449 Patent issued from U.S. Application No. 10/530,267 (“the ‘267 application”). During prosecution of the ‘267 application, the Examiner rejected the pending claims under 35 U.S.C. § 103(a) as being unpatentable over Lin in view of U.S. Pub. No. 2003/0182419 (“Barr”). *See* Ex. 1002, 53–59 (Office Action dated March 2, 2007). In response, Applicant made minor amendments to the pending claims and asserted that “Lin and Barr fail to disclose or suggest the features of claim 1–7 which include an Integrated circuit comprising a plurality of processing module (M; I; S; T) and a network (N; RN) ...” because “Lin describes a ‘cluster’ which is two or more separate computer systems” and “Barr describes a

telematics client framework and does not rectify the shortcomings of Lin as to this rejection.” *See* Ex. 1002 at 48–49.

86. Thereafter, the Examiner rejected the pending claims under 35 U.S.C. § 103(a) as being unpatentable over Emerson in view of U.S. Patent No. 6,769,046 (“Adams”). Ex. 1002 at 31–37. In response, the Applicant amended the pending claims and asserted that the amended claims are patentable over Emerson in view of Adams because:

Emerson does not disclose or suggest the features of claim 1 which include the resource manager determining whether the at least one connection with the communication channels is available based on the connection properties being available and the features of claim 7 which include the first module issuing a request for a connection with the second module to a communication manager, the resource manager determining whether a target connection with the desired connection properties is available, and the resource manager responding with the availability of the target connection to the communication manager.

...

Similar to Emerson, Adams also does not disclose or suggest the features of claim 1 which include the resource manager determining whether the at least one connection with the communication channels is available based on the connection properties being available and the features of claim 7 which include the first module issuing a request for a connection with the second module to a communication manager, the resource manager determining whether a target connection with the desired connection properties is available, and the resource manager responding with the availability of the target connection to the communication manager.

See Ex. 1002 at 20–27.

87. Following the applicant's amendment and response of November 20, 2007, the Examiner granted the '267 application.

88. To the extent the prosecution history is specifically relevant to my opinions, I will note it in the paragraphs below.

VIII. CLAIM CONSTRUCTION

89. I understand that the Petition states that "the Board does not need to construe any claim term for purposes of evaluating the prior art in this Petition" and does not propose any claim constructions. Petition 15. Similarly, Dr. Mowry does not offer any opinions on claim construction in his declaration.

90. I understand that the parties in the various co-pending district court litigations have proposed, and in some cases agreed to, certain claim constructions. I also understand that the court issued a claim construction order in the TI Litigation. My understanding of these claim construction issues is summarized below.

A. The “[first/second] module,” “communication manager” and “resource manager” are separate and distinct (Claims 2, 6, and 9-16)

91. The terms “[first/second] module,” “resource manager” and “communication manager” are recited in all of the challenged claims.¹ Independent claim 10 recites that the first module issues “a request for a connection with the second module to a communication manager.” Ex. 1001 at 19:33–35. Next, the communication manager “forward[s] the request to a resource manager.” *Id.* at 19:36–37. Then, the resource manager “determine[es] whether a target connection with the desired connection properties is available” and “respond[s] with the availability of the target connection to the communication manager.” Claims 2 and 6 contain similar limitations and expressly include both a “communication manager” and a “resource manager.” *Id.* at 19:38–20:2.

92. In my opinion, a POSITA would understand that the claimed “[first/second] module,” “communication manager,” and “resource manager” are separate and distinct components or elements of the claimed invention.

¹ I understand Patent Owner has statutorily disclaimed claims 1 and 3–5 of the ‘449 Patent and that these claims are not at issue in this proceeding.


93. At the outset, I note that the terms “first module,” “second module,” “communication manager,” and “resource manager” are separately recited in the claims.

94. The claims require specific interactions or communications between the modules. For example, claim 10 recites “messages between the modules being exchanged over connections via a network.” Claim 10 also recites “at least one of outgoing messages from the first module to the second module and return messages from the second module to the first module.” Based on this information exchange, a POSITA would recognize that the first module and the second module are separate and distinct. Applying the same logic, a POSITA would recognize that the first module is separate and distinct from the communications manager because claim 10 recites that the first module “issu[es] a request for a connection with the second module to a communication manager.” The fact that there is an information exchange from the first module to the communication manager reinforces that the first module and the communication manager are separate and distinct components of the claimed invention.

95. The claims of the ’449 Patent also require specific interactions between the “communication manager” and “resource manager.” For example, claim 10 requires that the first module sends a request to the communication manager. The communication manager then forwards this request to the resource

manager. This demonstrates that the communication manager and the resource manager are separate and distinct; otherwise, the communication manager would absurdly forward the request to itself. The resource manager determines whether a target connection with the desired connection properties is available and responds to the communication manager. This further demonstrates that the the communication manager and the resource manager are separate and distinct; otherwise, the resource manager would absurdly respond to itself. Construing the communication manager and the resource manager as a single component would vitiate the specific interactions and exchanged of information as required by the claims.

96. This understanding of the “[first/second] module,” “communication manager,” and “resource manager” is confirmed by the specification of the ‘449 Patent. FIG. 3 (reproduced below) illustrates separate and distinct “CM” (communication manager) and “RM” (resource manager) elements in the router network RN.



Ex. 1001 at FIG. 3

97. Further, the specification of the '449 Patent explains that “the communication manager CM and the resource manager RM are arranged in the router network RN, but they can also be arranged in one or some of the network interfaces NI.” Ex. 1001 at 10:3–6. The '449 Patent does not, however, disclose or suggest that the communication manager could be disposed in a processing module, such as the initiator module “I” shown in FIG. 3 of the '449 Patent.

98. Indeed, the '449 Patent confirms that the “communication manager” is a separate and distinct component from a first module or a second module because the '449 Patent extolls the decoupling of IP-module implementation from network implementation, teaching away from implementing the communication manager within an IP module (e.g., a processing module):

The objectives when defining network services are the following. First, the services should abstract from the network internals as much as possible. **This is a key ingredient in tackling the challenge of decoupling the computation from communication, which allows IPs**

(the computation part), and the interconnect (the communication part) to be designed independently from each other. As a consequence, the services are positioned at the transport layer in the OSI reference model, which is the first layer to be independent of the implementation of the network.

Ex. 1001 at 17:59–18:1 (emphasis added).

NoCs also decouple computation from communication, which is essential in managing the design of billion-transistor chips. NoCs achieve this decoupling because they are traditionally designed using protocol stacks, which provide **well-defined interfaces separating communication service usage from service implementation**

Id. at 1:44–49 (emphasis added).

A set of NoC services is defined that abstract from the network details. **Using these services in IP design decouples computation and communication.**

Id. at 15:61–63.

99. For at least these reasons, a POSITA would understand that the claimed “[first/second] module,” “communication manager,” and “resource manager” are separate and distinct components or elements of the claimed invention.

B. Samsung’s Litigation Construction for “communication manager” (claims 2, 6, and 9–16)

100. The term “communication manager” is recited in claims 2, 6, and 9–16 of the ‘449 Patent.

101. I understand that in the TI Litigation, the court held that “the term ‘communication manager’ does not connote sufficiently definite structure” and is

governed by 35 U.S.C. § 112, ¶ 6. Ex. 2022 at 40. Specifically, I understand the court construed “communication manager” as follows:

Function: forwarding a request to a resource manager.

Corresponding Structure: portion of the circuit configured to implement the scheme disclosed in the ‘449 Patent at 449 at 10:19–24, 10:41–54.

Id. at 42.

102. The ’449 Patent describes the communication manager as follows:

The communication manager CM request a connection with a set of properties between two modules from the resource manager RM, after receiving the request REQ from the active network interface ports ANIP, i.e. the properties of a connection, e.g. throughput, flow control, are to be requested when asking for a connection setup

...

After enquiring of the available connection properties, the resource manager RM reserves the available connection i.e. the required resources, and responds to the communication manager CM which connection is available, i.e. which connection properties or the required resources thereof are available for the desired channels. Optionally, the communication manager CM can accept the connection with the available properties, but on the other hand the communication manager CM may also refuse the offered connection if the available properties or the resources thereof are not acceptable. If the communication manager CM refuses the offered connection, it sends an error message to the initiator module I (i.e. via the ANIP) and requests the resource module RM to release the reservation of said connections.

Ex. 1001 at 10:19–24, 10:41–54.

103. I understand that Patent Owner and Petitioner have proposed the following constructions for “communication manager” in the Samsung Litigation:

Term	Petitioner's Construction	Patent Owner's Construction
"communication manager"	Subject to Section 112, ¶ 6 Function: forwarding a request to a resource manager Structure: portion of the circuit configured to implement the scheme disclosed in the '449 Patent at 10:19–24, 10:41–54 Otherwise, indefinite.	Plain and ordinary meaning If subject to Section 112, ¶ 6: Function: forward a request for a connection to the resource manager Structure: circuit

Ex. 2033 at Appendix 2. Petitioner's district court construction for

"communication manager" is the same as the construction adopted by the court in the TI Litigation.

104. I further understand that Qualcomm has proposed the same Section 112, ¶ 6 construction for "communication manager" that was adopted in the TI Litigation and proposed by Petitioner in the Samsung Litigation.

C. District Court Constructions

105. I understand that Petitioner agreed to several claim constructions in the Samsung Litigation, and that several claim constructions remain in dispute, as summarized below.

1. “Network” (Claims 2, 6, and 9–16)

106. I understand that in the Samsung Litigation, Petitioner and Patent Owner agreed that the term “network” should be construed as “multi-hop interconnection between modules separated by one or more network nodes.”

107. I also understand that the parties to the Qualcomm Litigation agreed to this same construction for “network.”

108. I further understand that the parties to the TI Litigation disputed the proper construction for “network” and that the court adopted Patent Owner’s construction, which is the same as the construction agreed to in the Samsung and Qualcomm Litigations.

109. The specification of the ’449 Patent describes a “network” on a chip as a radical change from prior on-chip interconnections:

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 with the prevalent existing interconnects (i.e., buses) where modules are directly connected.

Ex. 2022, 22; Ex. 1001 at 3:47–54. Accordingly, agree with the construction of “network” provided above, and that is the meaning I am applying herein.

2. “Integrated Circuit” (Claims 2, 6, and 9–16)

110. I understand that in the Samsung Litigation, Petitioner and Patent Owner agreed that the term “integrated circuit” should be construed as “interconnected circuitry on a chip.” Ex. 2033 at 5-6.

111. I understand that in the TI Litigation, the 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 Properties” (Claims 2, 6, and 9–16)

112. I understand that in the Samsung Litigation, Petitioner and Patent Owner agreed that the term “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), 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 5-6.

113. I understand that the parties in the TI and Qualcomm Litigations agreed to the same construction for “connection properties.” Ex. 2022 at 57; Ex. 2052 at 2.

4. Claim 10's Preamble

114. I understand that in the Samsung Litigation, Petitioner and Patent Owner agreed that the preamble of claim 10 should be construed as limiting. Ex. 2033 at Appendix 1. I also understand that the parties in the TI and Qualcomm Litigations agreed that the preamble of claim 10 is limiting. Ex. 2022 at 58; Ex. 2052 at 2.

5. “Independently Configurable” (Claim 10)

115. I understand that in the Qualcomm Litigation, the parties agreed that “independently configurable” should be construed as “able to be configured with connection properties that do not depend on those of another configuration.” Ex. 2052 at 1.

6. “Connection” (Claims 2, 6, and 9–16)

116. I understand that in the Qualcomm Litigation, Qualcomm proposed that “connection” should be construed as “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.” Ex. 2052 at 1. I understand that Patent Owner proposed that “connection” should be construed as “one or more communication channels.” *Id.*

7. “At Least one of a Switch and a Router” (Claim 11)

117. I understand that in the TI and Qualcomm Litigations, the parties agreed that “at least one of a switch and a router” should be construed as “at least one switch and at least one router.” Ex. 2022 at 57; Ex. 2052 at 2.

8. “Flow Control” (Claim and 13)

118. I understand that in the Samsung Litigation, the parties agreed that “flow control” should be construed as “controlling flow whereby 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.” Ex. 2033 at 5. I also understand that the parties agreed to the same construction in the TI and Qualcomm Litigations. Ex. 2022 at 58; Ex. 2052 at 1.

IX. ANALYSIS OF THE ’449 PATENT

A. Ground 1: Goossens2002 in View of Drake Does Not Render the Challenged Claims of the ’449 Patent Obvious

1. Independent Claim 10

119. Claim 10 recites “the first module issuing a request for a connection with the second module to a communication manager” and “the communication manager forwarding the request to a resource manager.” Ex. 1001 at 19:32–37. The relevant steps (and parts thereof) of claim 10 of the ’449 Patent are reproduced below:

the first module issuing a request for a connection with the second module to a communication manager, . . . ;

the communication manager forwarding the request to a resource manager;
the resource manager determining whether a target connection with the desired connection properties is available;
the resource manager responding with the availability of the target connection to the communication manager; and
the target connection between the first and second module being established based on the available properties of said communication channels of said connection.

Ex. 1001 at 19:30–20:6.

- a. **Goossens2002 in view of Drake Fails to Disclose “said connections comprises a set of communication channels each having a set of connection properties, any communication channel being independently configurable”**

120. Claim 10 of the '449 Patent is directed to exchanging messages in an integrated circuit where the messages are exchanged between modules over connections via a network. Claim 10 requires that “said connections comprises **a set of communication channels.**” Ex. 1001 at 19:24–25. Sec. IX.A of Dr.

Mowry’s Declaration alleges that Goossens2002 discloses a connection comprising a set of communication channels. Ex. 1003 at ¶ 100. I disagree.

121. Claim 10 of the '449 Patent further recites that “any communication channel being independently configurable.” Ex. 1001 at 19:24–27. Sec. IX.B of Dr. Mowry’s Declaration alleges that Goossens2002 discloses channels being independently configurable. Ex. 1003 at ¶¶ 103–106, 139. I disagree.

122. Dr. Mowry’s Declaration alleges that a POSITA would have understood that each of the alleged channels in Goossens2002 are independently configurable, or independently adjustable², because “rather than describing that each connection has certain properties, Goossens2002 describes that each channel has certain properties.” Ex. 1003 at ¶ 103. But Goossens2002 never says “each channel.” Indeed, the allegedly supporting passage of Goossens2002 quoted in Dr. Mowry’s Declaration refers only to “a channel,” not “each channel.” Ex. 1005 at 3 (“**A channel** can have a combined throughput and latency guarantee (GT), or be best effort (BE), with an unspecified finite latency upper bound.”). Goossens2002 further explains in the same paragraph that “[a] connection combines multiple channels.” Ex. 1005 at 3. Nothing about combining channels into a connection discloses or suggests that those channels are independently configurable. At best, Goossens2002 is silent as to whether channels are independently configurable, and Dr. Mowry’s allegation to the contrary is an improper hindsight reconstruction

² Claim 1 of the ’449 Patent recites “the connection properties of the different communication channels of said connection being **adjustable** independently,” whereas claim 10 of the ’449 Patent recites “any communication channel being independently **configurable**.”

according to my understanding of the obviousness analysis outlined in Sec. III.D above.

123. Paragraph 104 of Dr. Mowry’s Declaration looks to Goossens2003 for to support the allegation that a POSITA would understand Goossens2002’s connections to be channels, alleging that “Goossens2003 . . . says explicitly that the channels can be independently configured to have different properties, i.e., that they are independently adjustable.” However, Goossens2003 does not say this either explicitly or implicitly. Dr. Mowry’s allegation about Goossens2003 relies on a sentence, from a paragraph in Goossens2003 that describes “connection-based service provision [that] allows better differentiation of services,” in which request and response *communications*—i.e. connections—are independently configurable, not in which request and response *channels* of connections are independently configurable. *See* Ex. 1007 at 74. Indeed, Goossens2003 never once mentions or suggests channels of a connection, or that a connection can comprise a set of channels. The relevant paragraph of Goossens2003 is reproduced below, in which it evident that the communications referred to by Dr. Mowry refer to connections, not channels.

The ÆTHEREAL services include transactions of bus protocols, to ease migration from current bus-based systems to NOCs. But where in OCP [23] and VCI [22] **connections are used** only to relax transaction ordering, **we offer a more general connection property model**. Not only are more properties considered, but the **request and response communications can be independently configured** (e.g.

for flow control, and throughput guarantees), allowing more fine-grained resource management. **Connection-based service provision therefore allows better differentiation of services**, and allows users to make full use of NOC performance.

Ex. 1007 at 74 (emphasis added).

124. Goossens2003 further explains that such differentiated services are achieved via connections (not channels), that connections (not channels) are created with a requested service level, and that different connections (not channels) are created and closed independently.

The ÆTHEREAL NOC offers **differentiated services with connections**. A connection describes communication between one master and one or more slaves, with an associated service level, such as fifo transaction ordering, and maximum latency. **Connections must be created stating the requested service level**; this is the negotiation phase. The NOC either accepts or rejects the request for the connection. Connection acceptance may lead to resource reservations in the NOC, e.g. buffers or a link bandwidth percentage. After usage, closing a connection frees the resources. **Different connections are created and closed independently, possibly at different points in time.**”).

Ex. 1007 at 71 (emphasis added).

125. Goossens2003 further provides an example of such differentiated services achieved via connections (not channels), where a microprocessor requests a first connection with a first bandwidth and a second connection with a second bandwidth.

As an example, consider a microprocessor which **negotiates a connection to a memory with guaranteed bandwidth**, but without a constraint on transaction ordering. **The NOC that offers communication services will honour the request** if has sufficient resources available, and will reject it otherwise. **Later, the microprocessor may renegotiate its connection, e.g. to a higher**

bandwidth. The NOC will release the resources of the connection for use by other connections, when it is closed by the microprocessor.

Ex. 1007 at 67 (emphasis added).

126. In summary, Goossens2003 does not disclose or suggest: (1) that a connection can comprise a set of channels; or (2) that channels (which in my opinion are not disclosed in Goossens2003) can be independently configurable. Thus, I disagree with the allegation in paragraph 104 of Dr. Mowry's Declaration that Goossens2003 describes an understanding that a POSITA would have regarding the disclosure of Goossens2002. Goossens2002 does not disclose or suggest channels that are independently configurable, and nothing disclosed in Goossens2003 suggests otherwise.

127. Paragraphs 106–110 of Dr. Mowry's Declaration allege that connections that comprise independently configurable channels would be obvious over Goossens2002 in view of Drake. I disagree, and I provide an explanation in the following section because the reasons involve the “allocation request” disclosed in Drake. To summarize briefly, however, each allocation request in Drake concerns a single data stream. While data streams in Drake can be allocated with different QoS parameters, Drake does not disclose that a single allocation request can concern a set of data streams (e.g., a forward direction and a backward direction). This would be required to satisfy the limitation of claim 10 of the '449 patent that recites “a request for a connection . . . wherein **the request** comprises

desired connection properties **associated with the sets of communication**

channels,” where “said connections comprises a set of communication channels.”

Accordingly, Drake does not disclose or suggest independently configurable channels of a connection.

128. For at least the reasons described above, Goossens2002 in view of Drake does not disclose or suggest the limitation of claim 10 of the '449 Patent that recites “said connections comprises a set of communication channels each having a set of connection properties, any communication channel being independently configurable.”

- b. Goossens2002 in view of Drake Fails to Disclose “the first module issuing a request for a connection with the second module to a communication manager, wherein the request comprises desired connection properties associated with the sets of communication channels”**

129. Paragraph 106 of Dr. Mowry’s Declaration alleges that “Drake also describes two different channels: one in the ‘forward direction,’ for communications ‘from source station 10 to target station 29,’ and a second in the ‘backward direction,’ for communications from the target station 29 to the source station 10.’ Ex. 1006 (Drake) at 5:30–50.” However, this section of Drake describes the hierarchical database structure for QoS management of local-area network (LAN) facilities, where the alleged “forward direction” and “backward direction” correspond to “quality-of-service parameters that describe the service

expected from the network.” Ex. 1006 at 5:30–33. These are not channels of a connection that are independently configurable. In fact, the “AllocationRequest” message relied on in paragraph 107 of Dr. Mowry’s Declaration is used to “request an allocation of a QoS path **for a particular data stream**.” Ex. 1006 at 8:31–32 (emphasis added). Drake does not disclose or suggest that a data stream comprises channels, nor does Drake disclose or suggest that the requested QoS is for a set of channels, e.g., that the requested QoS is for a set of data streams.

130. To be clear, Drake teaches an “apparatus and a method to control multimedia data stream entry into a LAN, based on the quality of service requirements **for that particular data stream** and the quality of service available in the LAN.” Ex. 1006 at 2:14–18 (emphasis added); *see also* Ex. 1006 at 10:4–8 (“Field 55 contains a specific stream while field 56 contains a specification of the minimum quality of service which is acceptable **for this data stream**.”).

131. Drake explains that a request for a data stream consists of a request for QoS parameters for only that particular data stream. Drake 10:42–46 (“[I]n box 71, the source station 10 actually sends a request to the allocator 26 to determine if the QoS parameters required **for the data stream** is available on the transmission path between source station 10 and target station 29.”) Drake does not disclose or suggest that a request for a data stream comprises “desired connection properties [e.g., QoS parameters] associated with the **sets of communication channels**” of a

connection as required by claim 10 of the '449 Patent, where the channels are independently configurable. *See* Ex. 1001 at 19:34–35 (emphasis added). Drake's QoS allocation request ("AllocationRequest") includes a "QoSAcceptable" field, which indicates "[t]he quality of service acceptable **for this data stream.**" Drake 8:43–46 (emphasis added). Such acceptable QoS, if granted, applies *to the entire requested data stream*, not to sets of channels within the data stream. Ex. 1006 at 8:49–61 ("AllocationReply—Reply to an AllocationRequest. . . . QoS Granted. QoS parameters of the path **granted to data stream**, if granted, as described above." (emphasis added)).

132. Drake does not disclose or suggest that a data stream can be partitioned into channels, e.g., a forward channel and a backward channel as suggested in paragraph 107 of Dr. Mowry's Declaration. Ex. 1003 at ¶ 107 ("These [QoS] parameters include throughput and end-to-end delay (i.e., latency) that is separately specified for the forward direction channel and for the backward direction channel."). In Drake, a data stream, from a source to a target (e.g., in a forward direction), having first QoS parameters is achieved via a first allocation request from the source to the target, while a second data stream, from the target to the source (e.g., in a reverse direction), having second QoS parameters is achieved via *a second allocation request* from the target to the source. *See, e.g.,* Ex. 1006 at 4:27–30 ("More particularly, station 10 is the source of a desired multimedia

transmission and comprises an application 11 which requests a QoS connection to the destination or target station 29.”). Drake does not disclose a single allocation request comprising QoS parameters for both an alleged forward channel (e.g., forward data stream) and an alleged backward channel (e.g., backward data stream). Drake’s system would need to issue two requests: one for each data flow. In contrast, in claim 10 of the ’449 Patent, a single “request comprises desired connection properties associated with the **sets of communication channels**.” Ex. 1001 at 19:34–35 (emphasis added). In other words, the single request in claim 10 of the ’449 Patent comprises the connection properties (e.g., QoS parameters) for all channels of the connection (e.g., for both a request channel and a response channel).

133. Paragraph 107 of Dr. Mowry’s Declaration cites several forward and backward QoS parameters that are listed in Drake, including “Throughput—Forward Direction,” “Throughput—Backward Direction,” “End-to-End Delay—Forward Direction,” and “End-to-End —Forward Direction.” Paragraph 108 of Dr. Mowry’s Declaration contends that “[a] POSITA would have recognized that there would be no reason for Drake to break out, e.g., throughput for the forward channel and backward channel, unless the channels can be configured to guarantee different amounts of throughput.” I disagree. Drake explains that these parameters describe *the service expected from the network*, i.e., the services available in the

components in the multimedia transmission path, not the QoS parameters allocated to a particular data stream. Ex. 1006 at 5:30–34 (“While quality-of-service parameters that **describe the service expected from the network** can be specified in many different ways, for the purposes of this application, the following parameters are used: Throughput—Forward Direction. . . . Throughput—Backward Direction. . . . End-to-End Delay—Forward Direction. . . . End-to-End —Forward Direction.” (emphasis added)); *see also, e.g.*, Drake at 10:58–11:2 (“If more components remain on the list, box 83 is entered where the next component identification is retrieved from the list. In box 84 the data base 25 of FIG. 1 is searched for the identified component parameters. . . . If the component parameters are found, box 88 is entered to increment the available path delay parameter, initialized in box 81, by the amount of delay in the current component. The decision box 89 is then entered to determine if the path delay parameter accumulated so far is not greater than the maximum delay QoS parameter required for this multimedia file.”).

134. For at least the reasons described above, Goossens2002 in view of Drake does not disclose or suggest the limitation of claim 10 of the ’449 Patent that recites “the first module issuing a request for a connection with the second module to a communication manager, wherein the request comprises desired connection properties associated with the sets of communication channels.”

- c. **Goossens2002 in view of Drake Fails to Disclose “the first module issuing a request for a connection with the second module to a communication manager, . . . the communication manager forwarding the request to a resource manager, . . . the resource manager responding with the availability of the target connection to the communication manager”**

135. Claim 10 of the '449 Patent describes a communication manager and a resource manager, where the communication manager receives a request for a connection from a first module, the communication manager forward the request to the resource manager, and the resource manager responds with connection availability to the communication manager. The relevant steps (and parts thereof) of claim 10 of the '449 Patent are reproduced below:

the first module issuing a request for a connection with the second module to a communication manager, . . . ;
the communication manager forwarding the request to a resource manager;
the resource manager determining whether a target connection with the desired connection properties is available;
the resource manager responding with the availability of the target connection to the communication manager; and
the target connection between the first and second module being established based on the available properties of said communication channels of said connection.


Ex. 1001 at 19:30–20:6.

136. Dr. Mowry's Declaration concedes that Goossens2002 fails to disclose a communication manager as described in Claim 10 of the '449 Patent and alleges that Drake's "QoS requestor manager" discloses the communication

manager. Ex. 1003 at ¶ 118. Paragraph 91 of Dr. Mowry’s Declaration also alleges that Drake’s “QoS allocator” discloses a resource manager as recited in Claim 10 of the ’449 Patent.

137. First, Drake’s QoS requestor manager does not disclose or suggest the communication manager of the ’449 Patent, and Drake’s QoS allocator does not disclose or suggest the resource manager of the ’449 Patent, at least because Drake does not disclose or suggest “a request for a connection” as explained in IX.A.1.a above. For example, the communication manager forwards the connection request to the resource manager in claim 10 of the ’449 Patent. But, because the connection request is missing from Drake, Drake’s QoS requestor manager does not perform the forwarding of the request and Drake’s QoS allocator does not perform the receiving of the request.

138. Next, Dr. Mowry’s Declaration points to two different elements in Drake’s QoS management system as allegedly disclosing the communication manager of the ’449 Patent, as indicated in the annotated version of FIG. 1 of Drake that is provided in paragraph 121 of Dr. Mowry’s Declaration (reproduced below).



However, neither of these elements—the QoS requestor manager 12 in the source station 10, or the QoS requestor manager 23 in the QoS allocator 20—discloses the communication manager of the '449 Patent, as I explain below.

(i) Drake's QoS requestor manager 12

139. Referring to FIG. 1 of Drake, “station 10 is the source of a desired multimedia transmission.” Drake 4:26–35. The source station 10 includes a QoS requestor manager 12 and a QoS protocol machine 13, which together assemble a request for a QoS data stream. Ex. 1006 at 4:32–37 (“Source station 10 also includes a QoS requestor manager 12 and a QoS protocol machine 13. Together,

components 12 and 13 assemble a request to reserve a QoS connection for a multimedia data stream on that path through the LAN 17 between source station 10 and target station 29.”); *see also* Ex. 1006 at 14:2–7 (“In accordance with the present invention, each source station, like source station 10 in FIG. 1, must continually refresh its request for the QoS data stream allocation at predetermined intervals or the requested bandwidth will be de-allocated to make this bandwidth available for other connections.”). In other words, source station 10 is a source of a data stream, and also, station 10 issues, via its QoS requestor manager 12, a request for the data stream.

140. Paragraphs 118–119 of Dr. Mowry’s Declaration allege that the QoS requestor manager 12 of station 10 discloses the communication manager of the ’449 Patent. Paragraph 97 of Dr. Mowry’s Declaration additionally alleges that the source and target stations of Drake disclose the first and second modules of the ’449 Patent. Ex. 1003 at ¶ 97 (“Drake similarly describes “connections” between source and target stations (i.e., first and second processing modules).”).

141. Thus, according to Dr. Mowry’s Declaration, Drake’s station 10 discloses *both* the first module (e.g., an initiator module) *and* the communication manager that are described in claim 10 of the ’449 Patent. However, claim 10 of the ’449 Patent recites “**the first module issuing a request** for a connection with the second module **to a communication manager.**” Ex. 1001 at 19:32–33

(emphasis added). Accordingly, a POSITA would recognize that the first module and the communication manager described in claim 10 of the '449 Patent are separate and distinct elements because an element would not issue a request for a connection to itself. For the same reason, a POSITA would recognize that Drake's station 10, which is the source of a data stream, would not issue a request for the data stream to itself via its own QoS requestor manager 12.

142. FIG. 3 of the '449 Patent (reproduced below) shows separate and distinct "CM" (communication manager) and "RM" (resource manager) elements in the router network RN.




FIG.3

Further, the specification of the '449 Patent explains that "the communication manager CM and the resource manager RM are arranged in the router network RN, but they can also be arranged in one or some of the network interfaces NI." Ex. 1001 at 10:3–6. The '449 Patent does not disclose or suggest that the communication manager could be disposed in a processing module, such as the

initiator module “I” shown in FIG. 3 of the ’449 Patent. Indeed, the ’449 Patent teaches away from implementing the communication manager within a processing module by extolling the decoupling of IP modules from the network.

The objectives when defining network services are the following. First, the services should abstract from the network internals as much as possible. **This is a key ingredient in tackling the challenge of decoupling the computation from communication, which allows IPs (the computation part), and the interconnect (the communication part) to be designed independently from each other.** As a consequence, the services are positioned at the transport layer in the OSI reference model, which is the first layer to be independent of the implementation of the network.

Ex. 1001 at 17:59–18:1 (emphasis added).

NoCs also decouple computation from communication, which is essential in managing the design of billion-transistor chips. NoCs achieve this decoupling because they are traditionally designed using protocol stacks, which provide **well-defined interfaces separating communication service usage from service implementation**

Ex. 1001 at 1:44–49 (emphasis added).

A set of NoC services is defined that abstract from the network details. **Using these services in IP design decouples computation and communication.**

Ex. 1001 at 15:61–63. Goossens2002 similarly extolls decoupling computation from communication, emphasizing the importance of re-using IP blocks in systems on silicon while relegating the communication aspects to the elements of the on-chip network.

We believe in the intellectual-property block (IP) re-use model. This means that a SOS contains hundreds of IPs that are bought or

premaded. Assuming their correctness, their communication is then the problem. We propose a network on silicon (NOS) to implement the communication between these IPs.

Ex. 1005 at 1.

Although a NOS based on packet-switching has its costs (packetisation, routers, buffering, harder to give guarantees) it is essential for an IP re-use strategy because it enables compositional and scalable integration of the IPs. Moreover, a NOS infrastructure (such as routers and protocol stack) can be re-used and its design cost amortised over multiple designs.

Ex. 1005 at 2.

In a re-use model IPs that are bought or re-used cannot be adapted to each SOS they are used in.

Ex. 1005 at 2.

143. Accordingly, a POSITA would lack motivation to modify Goossens2002's IP blocks to implement Drake's QoS requestor manager 12 of the station 10.

144. For at least the above reasons, Drake's QoS requestor manager 12 does not disclose or suggest the communication manager described in claim 10 of the '449 Patent.

(ii) Drake's QoS requestor manager 23


145. Referring again to FIG. 1 of Drake, "[a]llocator station 20 also includes a QoS requestor manager 23 which, together with QoS protocol machine 24 receives, decodes and acts upon the QoS request from source station 10." Drake

5:5–8. Paragraph 121 of Dr. Mowry’s Declaration alleges that the QoS requestor 23, which is a part of the QoS allocator 20, discloses the communication manager of the ’449 Patent. Ex. 1003 at ¶ 121 (“Additionally, Drake further describes that a QoS requestor manager associated with the QoS allocator “receives, decodes and acts upon the QoS request from source station 10.” Ex-1006 (Drake) at 5:5–8. These QoS requestor managers are illustrated in Drake’s Figure 1 annotated below.”). Paragraphs 91–94 of Dr. Mowry’s Declaration additionally allege that Drake’s QoS allocator 20 discloses the resource manager of the ’449 Patent. Ex. 1003 at ¶ 91 (“Drake similarly discloses a ‘resource manager’ in the form of a “quality of service (QoS) allocator.”).

146. Thus, according to Dr. Mowry’s Declaration, Drake’s QoS allocator 20 discloses *both* the communication manager *and* the resource manager that are described in claim 10 of the ’449 Patent. However, claim 10 of the ’449 Patent recites “the communication manager forwarding the request to a resource manager.” Ex. 1001 at 19:36–37. Accordingly, a POSITA would recognize that the communication manager and the resource manager described in claim 10 of the ’449 Patent are separate and distinct elements because an element would not forward a request for a connection to itself. For the same reason, a POSITA would recognize that Drake’s QoS allocator 20, which is the recipient of the request for a

data stream, would not forward the request for the data stream to itself via its own QoS requestor manager 23.

147. FIG. 3 of the '449 Patent (reproduced below) shows separate and distinct “CM” (communication manager) and “RM” (resource manager) elements in the router network RN.



148. For at least the above reasons, Drake’s QoS requestor manager 23 does not disclose or suggest the communication manager described in claim 10 of the '449 Patent. Further, the specification of the '449 Patent explains that “the communication manager CM and the resource manager RM **are** arranged in the router network RN, but **they** can also be arranged in one or some of the network interfaces NI.” Ex. 1001 at 10:3–6 (emphasis added). The '449 Patent refers to the CM and RM as plural elements, not as a singular element, and these elements can be arranged in various parts of the NoC.

149. For at least the above reasons, Drake's QoS requestor manager 12 does not disclose or suggest the communication manager described in claim 10 of the '449 Patent.

150. In summary, Goossens2002 in view of Drake does not render obvious the limitations of claim 10 of the '449 Patent that recite "the first module issuing a request for a connection with the second module to a communication manager, . . . the communication manager forwarding the request to a resource manager, . . . the resource manager responding with the availability of the target connection to the communication manager."

d. A POSITA Would Not Combine Goossens2002 and Drake

151. In my opinion, a POSITA would not be motivated to combine Goossens2002 and Drake.

152. First, Goossens2002 concerns the design of systems on silicon (SoSs) that comprise numerous IP cores. Ex. 1005 at 1 ("We advocate a network on silicon (NOS) as a hardware architecture to implement communication between IP cores.). Goossens2002 explains that "[i]t will soon be possible to manufacture systems on silicon (SOSs) containing 10M gates and megabytes of embedded software. The complexity of these systems will be overwhelming." Ex. 1005 at 1. To this end, Goossens2002 describes "a network on silicon (NOS) [that] eases the

design of systems on silicon (SOSs) at the levels of hardware design and software programming.” Ex. 1005 at 3

153. In contrast, Drake concerns “[a] management system for local area networks [that] is attached to the network and receives requests from all other stations on the network.” Ex. 1006 at Abstract; *see also* Ex. 1006 at 2:25–28 (“In accordance with the illustrative embodiment of the present invention, multimedia data streams are managed in a local area network (LAN), or in a bridged LAN environment.”); *see also* Ex. 1006 at 2:57–59 (“A major advantage of the present invention is the ability to provide quality of service management in bridged LAN environments.”); *see also* Ex. 1006 at 17–22 (“Referring more particularly to FIG. 1, there is shown a detailed block diagram of a quality of service management system for local area networks comprising a prior art local area network (LAN) 17 LAN 17 may, for example, be a Token Ring Network.”).

154. The ’449 Patent explains that on-chip networks (e.g., NoCs, NoSs) differ drastically from off-chip networks (e.g., LANs). Ex. 1001 at 1:66–2:1 (“NoC's premises are different from off-chip networks, and, therefore, most of the network design choices must be reevaluated.”); *see also, e.g.*, Ex. 1001 at 2:1–5 (“On-chip networks have different properties (e.g., tighter link synchronization) and constraints (e.g., higher memory cost) leading to different design choices.”); Ex. 1001 at 2:16–17 (“For on-chip networks computation too comes at a relatively

high cost compared to off-chip networks.”); Ex. 1001 at 2:28–29 (“The number of wires and pins to connect network components is an order of magnitude larger on chip than off chip.”); Ex. 1001 at 2:43–46 (“A consequence of the tight on-chip resource constraints is that the network components (i.e., routers and network interfaces) must be fairly simple to minimize computation and memory requirements.”).

155. For at least the above reasons, a POSITA would not be motivated to modify the on-chip teachings of Goossens2002 with the off-chip teachings of Drake.

156. Next, as explained in Sec. IX.A.1.c.i above, paragraphs 118–119 of Dr. Mowry’s Declaration allege that Drake’s QoS requestor manager 12 of station 10 discloses the communication manager of the ’449 Patent, and paragraph 97 of Dr. Mowry’s Declaration alleges that Drake’s source station 10 disclose a processing module of the ’449 Patent. Paragraph 122 of Dr. Mowry’s Declaration asserts that POSITA would have found it obvious to look to Drake’s QoS requestor manager 12 within the station 10 to implement the communication manager of the ’449 Patent. But this is contrary to the teachings of Goossens2002, which explicitly emphasize the importance of re-using *premade* IP blocks in systems on silicon rather than custom-designing IP blocks to implement guaranteed services for the

communication network. Indeed, Goossens2002 teaches away from the above-described combination that is alleged in Dr. Mowry's Declaration.

We believe in the intellectual-property block (IP) re-use model. This means that a SOS contains hundreds of IPs that are bought or premade. Assuming their correctness, their communication is then the problem. We propose a network on silicon (NOS) to implement the communication between these IPs.

...

Guaranteed services are the focus of this paper: they serve as . .
. a requirement for the NOS hardware architecture.

Ex. 1005 at 1.

Although a NOS based on packet-switching has its costs (packetisation, routers, buffering, harder to give guarantees) it is essential for an IP re-use strategy because it enables compositional and scalable integration of the IPs. Moreover, a NOS infrastructure (such as routers and protocol stack) can be re-used and its design cost amortised over multiple designs.

Ex. 1005 at 2.

In a re-use model IPs that are bought or re-used cannot be adapted to each SOS they are used in. They will interact in many different ways (event-driven, data streaming, message passing, shared memory, etc.) [8]. Thus the interconnecting network has to be flexible, in terms of the services that it offers.

Ex. 1005 at 2.

157. For at least the above reasons, a POSITA would not be motivated to modify the IP blocks disclosed in Goossens2002 to implement the QoS requestor manager 12 of Drake.

158. Still further, claim 10 of the '449 Patent recites that a first processing module issues a request for a connection, and the specification teaches that “connections are assumed to be created, closed, or modified **infrequently**.” Ex. 1001 at 11:61–62 (emphasis added). However, Drake states that “each source station, like source station 10 in FIG. 1 [(e.g., a first processing module according to Dr. Mowry’s Declaration)], **must continually refresh its request for the QoS data stream allocation** at predetermined intervals or the requested bandwidth will be de-allocated to make this bandwidth available for other connections.” Ex. 1006 at 14:3–7 (emphasis added); *see also* Ex. 1006 at 3:8–9 (“QoS allocations are maintained only so long as a requestor continues to refresh the request.”). Further, as suggested in Goossens2002, there is a desire for efficiently implementing performance guarantees in a NoC. Ex. 1005, 2 (“Performance guarantees The way packets are buffered and scheduled in routers, and the effects on performance guarantees has been the subject of intense research. Fundamentally, sharing and guarantees are conflicting, and efficiently combining guaranteed traffic with best-effort traffic is hard.”). Thus, a POSITA would lack motivation to look to Drake at least because the '449 Patent teaches away from frequent connection requests while Drake’s QoS management system explicitly implements frequent data stream allocation requests.

2. Claims 2, 6, 9, and 11–16

159. Claims 2, 6, 9, and 11–16 are not obvious over Goossens2002 in view of Drake for at least the same reasons that independent claim 10 is not obvious over Goossens2002 in view of Drake.

B. Ground 2: Goossens2003 in view of Drake Does Not Render Claims 1–6 or 9–16 of the ‘449 Patent Obvious

1. Independent Claim 10

- a. Goossens2003 in view of Drake Fails to Disclose “said connections comprises a set of communication channels each having a set of connection properties, any communication channel being independently configurable”**

160. Claim 10 of the ’449 Patent is directed to exchanging messages in an integrated circuit where the messages are exchanged between modules over connections via a network. Claim 10 requires that “said connections comprises a set of communication channels.” Ex. 1001 at 19:24–25. Sec. IX.B of Dr. Mowry’s Declaration alleges that Goossens2003 discloses a connection comprising a set of communication channels, and these channels can be independently configured. Ex. 1003 at ¶ 178; *see also* Ex. 1003 at ¶ 104 (“Goossens2003 . . . says explicitly that the channels can be independently configured to have different properties, i.e., that they are independently adjustable.””). However, Goossens2003 does not say this either explicitly or implicitly. Dr. Mowry’s allegation relies on a sentence, from a paragraph in Goossens2003 that describes “connection-based

service provision [that] allows better differentiation of services,” in which request and response *communications*—i.e. connections—are independently configurable, not in which request and response *channels* of connections are independently configurable. *See* Ex. 1007 at 74. Indeed, Goossens2003 never once mentions or suggests channels of a connection, or that a connection can comprise a set of channels. The relevant paragraph of Goossens2003 is reproduced below, in which it is evident that the communications referred to by Dr. Mowry refer to connections, not channels.

The ÆTHEREAL services include transactions of bus protocols, to ease migration from current bus-based systems to NOCs. But where in OCP [23] and VCI [22] **connections are used** only to relax transaction ordering, **we offer a more general connection property model**. Not only are more properties considered, but the **request and response communications can be independently configured** (e.g. for flow control, and throughput guarantees), allowing more fine-grained resource management. **Connection-based service provision therefore allows better differentiation of services**, and allows users to make full use of NOC performance.

Ex. 1007 at 74 (emphases added).

161. Goossens2003 further explains that such differentiated services are achieved via connections (not channels), that connections (not channels) are created with a requested service level, and that different connections (not channels) are created and closed independently.

The ÆTHEREAL NOC offers **differentiated services with connections**. A connection describes communication between one master and one or more slaves, with an associated service level, such as fifo transaction ordering,

and maximum latency. **Connections must be created stating the requested service level;** this is the negotiation phase. The NOC either accepts or rejects the request for the connection. Connection acceptance may lead to resource reservations in the NOC, e.g. buffers or a link bandwidth percentage. After usage, closing a connection frees the resources. **Different connections are created and closed independently, possibly at different points in time.**”).

Ex. 1007 at 71 (emphases added).

162. Goossens2003 further provides an example of such differentiated services achieved via connections (not channels), where a microprocessor requests a first connection with a first bandwidth and a second connection with a second bandwidth.

As an example, consider a microprocessor which **negotiates a connection to a memory with guaranteed bandwidth,** but without a constraint on transaction ordering. **The NOC that offers communication services will honour the request** if has sufficient resources available, and will reject it otherwise. **Later, the microprocessor may renegotiate its connection, e.g. to a higher bandwidth.** The NOC will release the resources of the connection for use by other connections, when it is closed by the microprocessor.

Ex. 1007 at 67 (emphases added).

163. In summary, Goossens2003 does not disclose or suggest: (1) that a connection can comprise a set of channels; or (2) that channels (which in my opinion are not disclosed in Goossens2003) can be independently configurable.

164. Paragraph 179 of Dr. Mowry’s Declaration alleges that “Drake describes connections with two channels—the “forward direction” channel and the “backward direction” channel—for which “throughput” and “end-to-end delay”

(i.e., latency)—can be adjusted independently.” Ex. 1003 ¶ 179. As I explained in Sec. IX.A.1.b, Drake does not disclose or suggest: (1) that a connection can comprise a set of channels; or (2) that channels (which in my opinion are not disclosed in Drake) can be independently configurable. I disagree, and provide an explanation in the following section because the reasons involve the “allocation request” disclosed in Drake. To summarize briefly, however, each allocation request in Drake concerns a single data stream. While data streams in Drake can be allocated with different QoS parameters, Drake does not disclose that a single allocation request can concern a set of data streams (e.g., a forward direction and a backward direction). This would be required to satisfy the limitation of claim 10 of the ’449 patent that recites “a request for a connection . . . wherein **the request** comprises desired connection properties **associated with the sets of communication channels**,” where “said connections comprises a set of communication channels.” Accordingly, Drake does not disclose or suggest independently configurable channels of a connection.

165. For at least the reasons described above, Goossens2003 in view of Drake does not disclose or suggest the limitation of claim 10 of the ’449 Patent that recites “said connections comprises a set of communication channels each having a set of connection properties, any communication channel being independently configurable.”

- b. **Goossens2003 in view of Drake Fails to Disclose “the first module issuing a request for a connection with the second module to a communication manager, wherein the request comprises desired connection properties associated with the sets of communication channels”**

166. Paragraph 179 of Dr. Mowry’s Declaration refers to paragraph 106 thereof for the allegation that “Drake also describes two different channels: one in the ‘forward direction,’ for communications ‘from source station 10 to target station 29,’ and a second in the ‘backward direction,’ for communications from the target station 29 to the source station 10.’ Ex. 1006 (Drake) at 5:30–50.” Ex. 1003 at ¶ 106. However, this section of Drake describes the hierarchical database structure for QoS management of local-area network (LAN) facilities, where the alleged “forward direction” and “backward direction” correspond to “quality-of-service parameters that describe the service expected from the network.” Ex. 1006 at 5:30–33. These are not channels of a connection that are independently configurable. In fact, the “AllocationRequest” message relied on in paragraph 107 of Dr. Mowry’s Declaration is used to “request an allocation of a QoS path **for a particular data stream.**” Ex. 1006 at 8:31–32 (emphasis added). Drake does not disclose or suggest that a data stream comprises channels, nor does Drake disclose or suggest that the requested QoS is for a set of channels, e.g., that the requested QoS is for a set of data streams.

167. To be clear, Drake teaches an “apparatus and a method to control multimedia data stream entry into a LAN, based on the quality of service requirements **for that particular data stream** and the quality of service available in the LAN.” Ex. 1006 at 2:14–18 (emphasis added); *see also* Ex. 1006 at 10:4–8 (“Field 55 contains a specific stream while field 56 contains a specification of the minimum quality of service which is acceptable **for this data stream.**”).

168. Drake explains that a request for a data stream consists of a request for QoS parameters for only that particular data stream. Ex. 1006 at 10:42–46 (“[I]n box 71, the source station 10 actually sends a request to the allocator 26 to determine if the QoS parameters required **for the data stream** is available on the transmission path between source station 10 and target station 29.”) Drake does not disclose or suggest that a request for a data stream comprises “desired connection properties [e.g., QoS parameters] associated with the **sets of communication channels**” of a connection as required by claim 10 of the ’449 Patent, where the channels are independently configurable. *See* Ex. 1001 at 19:34–35 (emphasis added). Drake’s QoS allocation request (“AllocationRequest”) includes a “QoSAcceptable” field, which indicates “[t]he quality of service acceptable **for this data stream.**” Ex. 1006 at 8:43–46 (emphasis added). Such acceptable QoS, if granted, applies *to the entire requested data stream*, not to sets of channels within the data stream. Ex. 1006 at 8:49–61 (“AllocationReply—Reply to an

AllocationRequest. . . QoS Granted. QoS parameters of the path **granted to data stream**, if granted, as described above.” (emphasis added)).

169. Drake does not disclose or suggest that a data stream can be partitioned into channels, e.g., a forward channel and a backward channel as suggested in paragraph 107 of Dr. Mowry’s Declaration. Ex. 1003 at ¶ 107 (“These [QoS] parameters include throughput and end-to-end delay (i.e., latency) that is separately specified for the forward direction channel and for the backward direction channel.”). In Drake, a data stream, from a source to a target (e.g., in a forward direction), having first QoS parameters is achieved via a first allocation request from the source to the target, while a second data stream, from the target to the source (e.g., in a reverse direction), having second QoS parameters is achieved via *a second allocation request* from the target to the source. *See, e.g.*, Ex. 1006 at 4:27–30 (“More particularly, station 10 is the source of a desired multimedia transmission and comprises an application 11 which requests a QoS connection to the destination or target station 29.”). Drake does not disclose a single allocation request comprising QoS parameters for both an alleged forward channel (e.g., forward data stream) and an alleged backward channel (e.g., backward data stream). Drake’s system would need to issue two requests: one for each data flow. In contrast, in claim 10 of the ’449 Patent, a single “request comprises desired connection properties associated with the **sets of communication channels**.” ’449

Patent 19:34–35 (emphasis added). In other words, the single request in claim 10 of the '449 Patent comprises the connection properties (e.g., QoS parameters) for all channels of the connection (e.g., for both a request channel and a response channel).

170. Paragraph 107 of Dr. Mowry's Declaration cites several forward and backward QoS parameters that are listed in Drake, including "Throughput—Forward Direction," "Throughput—Backward Direction," "End-to-End Delay—Forward Direction," and "End-to-End —Forward Direction." Paragraph 108 of Dr. Mowry's Declaration contends that "[a] POSITA would have recognized that there would be no reason for Drake to break out, e.g., throughput for the forward channel and backward channel, unless the channels can be configured to guarantee different amounts of throughput." I disagree. Drake explains that these parameters describe *the service expected from the network*, i.e., the services available in the components in the multimedia transmission path, not the QoS parameters allocated to a particular data stream. Ex. 1006 at 5:30–34 ("While quality-of-service parameters that **describe the service expected from the network** can be specified in many different ways, for the purposes of this application, the following parameters are used: Throughput—Forward Direction. . . . Throughput—Backward Direction. . . . End-to-End Delay—Forward Direction. . . . End-to-End —Forward Direction." (emphasis added)); *see also, e.g.*, Drake at 10:58–11:2 ("If more

components remain on the list, box 83 is entered where the next component identification is retrieved from the list. In box 84 the data base 25 of FIG. 1 is searched for the identified component parameters. . . . If the component parameters are found, box 88 is entered to increment the available path delay parameter, initialized in box 81, by the amount of delay in the current component. The decision box 89 is then entered to determine if the path delay parameter accumulated so far is not greater than the maximum delay QoS parameter required for this multimedia file.”).

171. For at least the reasons described above, Goossens2003 in view of Drake does not disclose or suggest the limitation of claim 10 of the ’449 Patent that recites “the first module issuing a request for a connection with the second module to a communication manager, wherein the request comprises desired connection properties associated with the sets of communication channels.”

- c. **Goossens2002 in view of Drake Fails to Disclose “the first module issuing a request for a connection with the second module to a communication manager, . . . the communication manager forwarding the request to a resource manager, . . . the resource manager responding with the availability of the target connection to the communication manager”**

172. Paragraph 173 of Dr. Mowry’s Declaration concedes that Goossens2003 “does not describe a component or mechanism for tracking resource usage within the NOC or for determining the availability of the resources required

for a requested connection,” and looks to Drake for disclosure of the resource manager described in the ’449 Patent. As I explained in Sec. IX.A.1.c, Drake does not disclose either the resource manager or the communication manager described in the ’449 Patent. That explanation is reproduced below, but considering Goossens2003 instead of Goossens2002.

173. Claim 10 of the ’449 Patent describes a communication manager and a resource manager, where the communication manager receives a request for a connection from a first module, the communication manager forward the request to the resource manager, and the resource manager responds with connection availability to the communication manager. The relevant steps (and parts thereof) of claim 10 of the ’449 Patent are reproduced below:

- the first module issuing a request for a connection with the second module to a communication manager, . . . ;
- the communication manager forwarding the request to a resource manager;
- the resource manager determining whether a target connection with the desired connection properties is available;
- the resource manager responding with the availability of the target connection to the communication manager; and
- the target connection between the first and second module being established based on the available properties of said communication channels of said connection.


Ex. 1001 at 19:30–20:6.

174. Dr. Mowry’s Declaration alleges that Drake’s “QoS requestor manager” discloses the communication manager. Ex. 1003 at ¶ 118. Paragraph 91

of Dr. Mowry's Declaration also alleges that Drake's "QoS allocator" discloses a resource manager as recited in Claim 10 of the '449 Patent

175. First, Drake's QoS requestor manager does not disclose or suggest the communication manager of the '449 Patent, and Drake's QoS allocator does not disclose or suggest the resource manager of the '449 Patent, at least because Drake does not disclose or suggest "a request for a connection" as explained in IX.A.1.a above. For example, the communication manager forwards the connection request to the resource manager in claim 10 of the '449 Patent. But, because the connection request is missing from Drake, Drake's QoS requestor manager does not perform the forwarding of the request and Drake's QoS allocator does not perform the receiving of the request.

176. Next, Dr. Mowry's Declaration points to two different elements in Drake's QoS management system as allegedly disclosing the communication manager of the '449 Patent, as indicated in the annotated version of FIG. 1 of Drake that is provided in paragraph 121 of Dr. Mowry's Declaration (reproduced below).



However, neither of these elements—the QoS requestor manager 12 in the source station 10, or the QoS requestor manager 23 in the QoS allocator 20—discloses the communication manager of the '449 Patent, as I explain below.

(i) Drake's QoS requestor manager 12

177. Referring to FIG. 1 of Drake, “station 10 is the source of a desired multimedia transmission.” Ex. 1006 at 4:26–35. The source station 10 includes a QoS requestor manager 12 and a QoS protocol machine 13, which together assemble a request for a QoS data stream. Ex. 1006 at 4:32–37 (“Source station 10 also includes a QoS requestor manager 12 and a QoS protocol machine 13.


Together, components 12 and 13 assemble a request to reserve a QoS connection for a multi-media data stream on that path through the LAN 17 between source station 10 and target station 29.”); *see also* Ex. 1006 at 14:2–7 (“In accordance with the present invention, each source station, like source station 10 in FIG. 1, must continually refresh its request for the QoS data stream allocation at predetermined intervals or the requested bandwidth will be de-allocated to make this bandwidth available for other connections.”). In other words, source station 10 is a source of a data stream, and also, station 10 issues, via its QoS requestor manager 12, a request for the data stream.

178. Paragraphs 118–119 of Dr. Mowry’s Declaration allege that the QoS requestor manager 12 of station 10 discloses the communication manager of the ’449 Patent. Paragraph 97 of Dr. Mowry’s Declaration additionally alleges that the source and target stations of Drake disclose the first and second modules of the ’449 Patent. Ex. 1003 at ¶ 97 (“Drake similarly describes “connections” between source and target stations (i.e., first and second processing modules).”).

179. Thus, according to Dr. Mowry’s Declaration, Drake’s station 10 discloses *both* the first module (e.g., an initiator module) *and* the communication manager that are described in claim 10 of the ’449 Patent. However, claim 10 of the ’449 Patent recites “**the first module issuing a request** for a connection with the second module **to a communication manager.**” Ex. 1001 at 19:32–33

(emphasis added). Accordingly, a POSITA would recognize that the first module and the communication manager described in claim 10 of the '449 Patent are separate and distinct elements because an element would not issue a request for a connection to itself. For the same reason, a POSITA would recognize that Drake's station 10, which is the source of a data stream, would not issue a request for the data stream to itself via its own QoS requestor manager 12.

180. FIG. 3 of the '449 Patent (reproduced below) shows separate and distinct "CM" (communication manager) and "RM" (resource manager) elements in the router network RN.



Further, the specification of the '449 Patent explains that "the communication manager CM and the resource manager RM are arranged in the router network RN, but they can also be arranged in one or some of the network interfaces NI." Ex. 100 at 10:3–6. The '449 Patent does not disclose or suggest that the communication manager could be disposed in a processing module, such as the

initiator module “I” shown in FIG. 3 of the ’449 Patent. Indeed, the ’449 Patent teaches away from implementing the communication manager within a processing module by extolling the decoupling of IP modules from the network.

The objectives when defining network services are the following. First, the services should abstract from the network internals as much as possible. **This is a key ingredient in tackling the challenge of decoupling the computation from communication, which allows IPs (the computation part), and the interconnect (the communication part) to be designed independently from each other.** As a consequence, the services are positioned at the transport layer in the OSI reference model, which is the first layer to be independent of the implementation of the network.

’449 Patent 17:59–18:1 (emphasis added).

NoCs also decouple computation from communication, which is essential in managing the design of billion-transistor chips. NoCs achieve this decoupling because they are traditionally designed using protocol stacks, which provide **well-defined interfaces separating communication service usage from service implementation.**

’449 Patent 1:44–49 (emphasis added).

A set of NoC services is defined that abstract from the network details. **Using these services in IP design decouples computation and communication.**

Ex. 1001 at 15:61–63. Goossens2003 similarly extolls decoupling computation from communication, emphasizing the importance of re-using IP blocks in systems on silicon while relegating the communication aspects to the elements of the on-chip network.

System-on-chip designers use networks on chip (NOC) to solve deep submicron problems, and to divide global problems into local,

decoupled problems. NOCs provide services through protocol stacks, and introducing guaranteed services enables IP re-use and platform-based design.

Ex. 1007 at 61.

This trend commenced some time ago with IP-based re-use, platform-based design [5]. The interest in decoupling and composing local solutions can be observed in hierarchical lay-out and wiring, local clocking strategies such as GALS (globally asynchronous, locally synchronous) [6], software-programmable fixed IPs or tiles [7, 8], and synthesisable [sic] tiles for dedicated silicon [9] or FPGAs [10, 11].

Ex. 1007 at 64 (brackets for references in original).

IP re-use and platform-based design aim to re-use applications and architectures by decoupling them. The use of NOCs and their associated network protocol stacks are a way to achieve this (cf. Section 1.5). Finally, it is our tenet that QOS, in particular guaranteed services, allow us to move towards globally predictable, locally predictable methods and solutions for SOC design. This, in combination with NOCs solves a need of recent architectures and design methods.

Ex. 1007 at 65.

The aim of platform-based design [5, 18] is to reduce the cost of system design through re-use of applications and architectures. A platform decouples applications and system architectures, by defining a template architecture and programming model. In other words, by limiting the freedom in which an application can be implemented on an architecture, the interdependence of application and architecture is concentrated and reduced (Figure 4.2(a)).

Ex. 1007 at 65.

NOCs help IP re-use and platform-based design to combine IPs, and NOC services aid in decoupling applications and architectures.

Ex. 1007 at 66.

The combination of NOCs and QoS is natural, through network protocol stacks, and beneficial for several reasons. It enables IP re-use and platform-based design to decouple applications and architectures.

Ex. 1007 at 66.

181. Accordingly, a POSITA would lack motivation to modify Goossens2003's IP blocks to implement Drake's QoS requestor manager 12 of the station 10.

182. For at least the above reasons, Drake's QoS requestor manager 12 does not disclose or suggest the communication manager described in claim 10 of the '449 Patent.


(ii) Drake's QoS requestor manager 23

183. Referring again to FIG. 1 of Drake, "[a]llocator station 20 also includes a QoS requestor manager 23 which, together with QoS protocol machine 24 receives, decodes and acts upon the QoS request from source station 10." Ex. 1006 at 5:5–8. Paragraph 121 of Dr. Mowry's Declaration alleges that the QoS requestor 23, which is a part of the QoS allocator 20, discloses the communication manager of the '449 Patent. Ex. 1003 at ¶ 121 ("Additionally, Drake further describes that a QoS requestor manager associated with the QoS allocator "receives, decodes and acts upon the QoS request from source station 10." Ex-1006 (Drake) at 5:5–8. These QoS requestor managers are illustrated in Drake's Figure 1 annotated below."). Paragraphs 91–94 of Dr. Mowry's Declaration

additionally allege that Drake's QoS allocator 20 discloses the resource manager of the '449 Patent. Ex. 1003 at ¶ 91 ("Drake similarly discloses a 'resource manager' in the form of a "quality of service (QoS) allocator.").

184. Thus, according to Dr. Mowry's Declaration, Drake's QoS allocator 20 discloses *both* the communication manager *and* the resource manager that are described in claim 10 of the '449 Patent. However, claim 10 of the '449 Patent recites "the communication manager forwarding the request to a resource manager." Ex. 1001 at 19:36–37. Accordingly, a POSITA would recognize that the communication manager and the resource manager described in claim 10 of the '449 Patent are separate and distinct elements because an element would not forward a request for a connection to itself. For the same reason, a POSITA would recognize that Drake's QoS allocator 20, which is the recipient of the request for a data stream, would not forward the request for the data stream to itself via its own QoS requestor manager 23.

185. FIG. 3 of the '449 Patent (reproduced below) shows separate and distinct "CM" (communication manager) and "RM" (resource manager) elements in the router network RN.



186. For at least the above reasons, Drake’s QoS requestor manager 23 does not disclose or suggest the communication manager described in claim 10 of the ’449 Patent. Further, the specification of the ’449 Patent explains that “the communication manager CM and the resource manager RM are arranged in the router network RN, but they can also be arranged in one or some of the network interfaces NI.” Ex. 1001 at 10:3–6 (emphasis added). The ’449 Patent refers to the CM and RM as plural elements, not as a singular element, and these elements can be arranged in various parts of the NoC.

187. For at least the above reasons, Drake’s QoS requestor manager 12 does not disclose or suggest the communication manager described in claim 10 of the ’449 Patent.

188. In summary, Goossens2003 in view of Drake does not render obvious the limitations of claim 10 of the ’449 Patent that recite “the first module issuing a request for a connection with the second module to a communication manager, . . .

the communication manager forwarding the request to a resource manager, . . . the resource manager responding with the availability of the target connection to the communication manager.”

d. A POSITA Would Not Combine Goossens2003 and Drake

189. In my opinion, a POSITA would not be motivated to combine Goossens2002 and Drake.

190. First, Goossens2003 concerns the design of systems on silicon (SoSs) that enables IP re-use. Ex. 1005 at 61 (“System-on-chip designers use networks on chip (NOC) to solve deep submicron problems, and to divide global problems into local, decoupled problems. NOCs provide services through protocol stacks, and introducing guaranteed services enables IP re-use and platform-based design. . . . The ÆTHEREAL NOC is an example of this approach, and forms the basis of a QOS-based design style, as advocated in this chapter.”). Goossens2003 explains that “IP re-use and platform-based design aim to re-use applications and architectures by decoupling them. The use of NOCs and their associated network protocol stacks are a way to achieve this.” Ex. 1005 at 65; *see also* Ex. 1005 at 66 (“NOCs help IP re-use and platform-based design to combine IPs, and NOC services aid in decoupling applications and architectures.”); Ex. 1005 at 66 (“The combination of NOCs and QOS . . . enables IP re-use and platform-based design to decouple applications and architectures”).

191. In contrast, Drake concerns “[a] management system for local area networks [that] is attached to the network and receives requests from all other stations on the network.” Drake at Abstract; *see also* Ex. 1006 at 2:25–28 (“In accordance with the illustrative embodiment of the present invention, multimedia data streams are managed in a local area network (LAN), or in a bridged LAN environment.”); *see also* Ex. 1006 at 2:57–59 (“A major advantage of the present invention is the ability to provide quality of service management in bridged LAN environments.”); *see also* Ex. 1006 at 4:17–22 (“Referring more particularly to FIG. 1, there is shown a detailed block diagram of a quality of service management system for local area networks comprising a prior art local area network (LAN) 17 LAN 17 may, for example, be a Token Ring Network.”).

192. The ’449 Patent explains that on-chip networks (e.g., NoCs, NoSs) differ drastically from off-chip networks (e.g., LANs). Ex. 1001 at 1:66–2:1 (“NoC's premises are different from off-chip networks, and, therefore, most of the network design choices must be reevaluated.”); *see also, e.g.*, Ex. 1001 at 2:1–5 (“On-chip networks have different properties (e.g., tighter link synchronization) and constraints (e.g., higher memory cost) leading to different design choices.”); Ex. 1001 at 2:16–17 (“For on-chip networks computation too comes at a relatively high cost compared to off-chip networks.”); Ex. 1001 at 2:28–29 (“The number of wires and pins to connect network components is an order of magnitude larger on

chip than off chip.”); Ex. 1001 at 2:43–46 (“A consequence of the tight on-chip resource constraints is that the network components (i.e., routers and network interfaces) must be fairly simple to minimize computation and memory requirements.”).

193. For at least the above reasons, a POSITA would not be motivated to modify the on-chip teachings of Goossens2003 with the off-chip teachings of Drake.

194. Next, as explained in Sec. IX.B.1.c.i above, paragraphs 118–119 of Dr. Mowry’s Declaration allege that Drake’s QoS requestor manager 12 of station 10 discloses the communication manager of the ’449 Patent, and paragraph 97 of Dr. Mowry’s Declaration alleges that Drake’s source station 10 disclose a processing module of the ’449 Patent. Paragraph 122 of Dr. Mowry’s Declaration asserts that POSITA would have found it obvious to look to Drake’s QoS requestor manager 12 within the station 10 to implement the communication manager of the ’449 Patent. But this is contrary to the teachings of Goossens2003, which explicitly emphasize the importance of re-using IP blocks in systems on chip rather than custom-designing IP blocks to implement guaranteed services for the communication network. Indeed, Goossens2003 teaches away from the above-described combination that is alleged in Dr. Mowry’s Declaration.

System-on-chip designers use networks on chip (NOC) to solve deep submicron problems, and to divide global problems into local,

decoupled problems. NOCs provide services through protocol stacks, and introducing guaranteed services enables IP re-use and platform-based design.

Ex. 1007 at 61.

This trend commenced some time ago with IP-based re-use, platform-based design [5]. The interest in decoupling and composing local solutions can be observed in hierarchical lay-out and wiring, local clocking strategies such as GALS (globally asynchronous, locally synchronous) [6], software-programmable fixed IPs or tiles [7, 8], and synthesisable [sic] tiles for dedicated silicon [9] or FPGAs [10, 11].

Ex. 1007 at 64 (brackets for references in original).

IP re-use and platform-based design aim to re-use applications and architectures by decoupling them. The use of NOCs and their associated network protocol stacks are a way to achieve this (cf. Section 1.5). Finally, it is our tenet that QOS, in particular guaranteed services, allow us to move towards globally predictable, locally predictable methods and solutions for SOC design. This, in combination with NOCs solves a need of recent architectures and design methods.

Ex. 1007 at 65.

The aim of platform-based design [5, 18] is to reduce the cost of system design through re-use of applications and architectures. A platform decouples applications and system architectures, by defining a template architecture and programming model. In other words, by limiting the freedom in which an application can be implemented on an architecture, the interdependence of application and architecture is concentrated and reduced (Figure 4.2(a)).

Ex. 1007 at 65.

NOCs help IP re-use and platform-based design to combine IPs, and NOC services aid in decoupling applications and architectures.

Ex. 1007 at 66.

The combination of NOCs and QoS is natural, through network protocol stacks, and beneficial for several reasons. It enables IP re-use and platform-based design to decouple applications and architectures.

Ex. 1007 at 66.

195. For at least the above reasons, a POSITA would not be motivated to modify the IP blocks disclosed in Goossens2003 to implement the QoS requestor manager 12 of Drake.

196. Still further, claim 10 of the '449 Patent recites that a first processing module issues a request for a connection, and the specification teaches that “connections are assumed to be created, closed, or modified **infrequently**.” Ex. 1001 at 11:61–62 (emphasis added). However, Drake states that “each source station, like source station 10 in FIG. 1 [(e.g., a first processing module according to Dr. Mowry’s Declaration)], **must continually refresh its request for the QoS data stream allocation** at predetermined intervals or the requested bandwidth will be de-allocated to make this bandwidth available for other connections.” Ex. 1006 at 14:3–7 (emphasis added); *see also* Ex. 1006 at 3:8–9 (“QoS allocations are maintained only so long as a requestor continues to refresh the request.”). Thus, a POSITA would lack motivation to look to Drake at least because the '449 Patent teaches away from frequent connection requests while Drake’s QoS management system explicitly implements frequent data stream allocation requests.

2. Claims 2, 6, 9, and 11–16

197. Claims 2, 6, 9, and 11–16 are not obvious over Goossens2003 in view of Drake for at least the same reasons that independent claim 10 is not obvious over Goossens2003 in view of Drake.

X. CONCLUSION

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

A handwritten signature in black ink, appearing to read "Erik Chmelar". The signature is fluid and cursive, with the first name "Erik" and last name "Chmelar" clearly distinguishable.

Date: March 28, 2024

Dr. Erik Chmelar