

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

QUALCOMM INCORPORATED,
Petitioners

v.

NETWORK SYSTEM TECHNOLOGIES, LLC,
Patent Owner

Case IPR2024-_____

U.S. Patent No. 7,373,449

Declaration of Todd Mowry, Ph.D.

TABLE OF CONTENTS

LIST OF EXHIBITS	V
I. BACKGROUND AND QUALIFICATIONS.....	9
II. MATERIALS AND OTHER INFORMATION CONSIDERED.....	12
III. UNDERSTANDING OF PATENT LAW	13
A. The Level of Ordinary Skill in the Art.....	13
B. Obviousness	14
IV. SUMMARY OF OPINIONS	16
V. OVERVIEW OF THE TECHNOLOGY	16
A. Systems-on-Chip (SoC)	16
B. Networks-on-Chip (NoC)	18
C. Connections.....	19
D. Admission Control	24
VI. THE '449 PATENT	25
A. Claims	26
B. Summary of the Specification.....	28
C. Summary of the Prosecution History	30
VII. LEVEL OF ORDINARY SKILL IN THE ART.....	32
VIII. OVERVIEW OF THE PRIOR ART	34
A. Overview of Goossens2002	34
B. Overview of Drake.....	35
C. Overview of Goossens2003	37
IX. SPECIFIC GROUNDS FOR PETITION.....	38

A.	Ground I: Claims 1–6 and 9–16 Are Rendered Obvious by Goossens2002 in View of Drake	38
1.	Claim 1.....	45
2.	Claim 2: “Integrated circuit according to claim 1, further comprising: at least one communication manager (CM) for managing the communication between different modules”	66
3.	Claim 3: “The Integrated circuit of claim 1, further comprising at least one of a switch and a router”	69
4.	Claim 4: “The Integrated circuit of claim 1, further comprising a chip, wherein the processing modules and the network are disposed on said chip”	70
5.	Claim 5: “The Integrated circuit of claim 1, wherein the connection properties comprise at least one of data transport ordering, flow control, throughput, latency, and lossiness”	70
6.	Claim 6.....	71
7.	Claim 9: “Integrated circuit according to claim 2, further comprising: at least one network interface means (NI), associated to each of said modules, for managing the communication between said modules and said network (N)”	77
8.	Claim 10.....	77
9.	Claim 11: “The method of claim 10, wherein the network manages traffic utilizing at least one of a switch and a router”	81
10.	Claim 12: “The method of claim 10, wherein the processing modules and the network are disposed on a chip”	81
11.	Claim 13: The method of claim 10, wherein the connection properties comprise at least one of data transport ordering, flow control, throughput, latency, and lossiness”	81
12.	Claim 14: “The method of claim 10, further comprising the resource manager determining whether the target connection with the desired connection properties is available based on reading of a centralized or distributed property table comprising properties associated with the network”	82

13.	Claim 15: “The method of claim 10, further comprising the resource manager reserving the target connection after determining whether the target connection with the desired connection properties is available”	84
14.	Claim 16: “The method of claim 10, wherein the connection properties comprise at least one of throughput, latency and jitter”	84
B.	Ground II: Claims 1-6 and 9-16 Are Rendered Obvious by Goossens2003 in view of Drake	85
1.	Claim 1.....	91
2.	Claim 2: “Integrated circuit according to claim 1, further comprising: at least one communication manager (CM) for managing the communication between different modules”	101
3.	Claim 3: “The Integrated circuit of claim 1, further comprising at least one of a switch and a router”	102
4.	Claim 4: “The Integrated circuit of claim 1, further comprising a chip, wherein the processing modules and the network are disposed on said chip”	104
5.	Claim 5: “The Integrated circuit of claim 1, wherein the connection properties comprise at least one of data transport ordering, flow control, throughput, latency, and lossiness”	104
6.	Claim 6.....	105
7.	Claim 9: “Integrated circuit according to claim 2, further comprising: at least one network interface means (NI), associated to each of said modules, for managing the communication between said modules and said network (N)”	108
8.	Claim 10.....	109
9.	Claim 11: “The method of claim 10, wherein the network manages traffic utilizing at least one of a switch and a router”	112
10.	Claim 12: “The method of claim 10, wherein the processing modules and the network are disposed on a chip”	112
11.	Claim 13: The method of claim 10, wherein the connection properties comprise at least one of data	

	transport ordering, flow control, throughput, latency, and lossiness”	113
12.	Claim 14: “The method of claim 10, further comprising the resource manager determining whether the target connection with the desired connection properties is available based on reading of a centralized or distributed property table comprising properties associated with the network”	113
13.	Claim 15: “The method of claim 10, further comprising the resource manager reserving the target connection after determining whether the target connection with the desired connection properties is available”	114
14.	Claim 16: “The method of claim 10, wherein the connection properties comprise at least one of throughput, latency and jitter”	114
X.	SECONDARY CONSIDERATIONS	115
XI.	CONCLUSION	116

LIST OF EXHIBITS

Exhibit No.	Description
1001	U.S. Patent No. 7,373,449 (the “’449 Patent”)
1002	File History of U.S. Pat. No. 7,373,449
1003	Declaration of Todd Mowry Ph.D. in Support of Petition for <i>Inter Partes</i> Review of U.S. Patent No. 7,373,449
1004	<i>Curriculum Vitae</i> of Todd Mowry Ph.D.
1005	K. Goossens, et al., <i>Networks on Silicon: Combining Best-Effort and Guaranteed Services</i> (“Goossens2002”)
1006	U.S. Patent No. 5,461,611 (“Drake”)
1007	K. Goossens, et al., <i>Guaranteeing the Quality of Services in Networks on Chip</i> (“Goossens2003”)
1008	Ahmed Amine Jerraya & Wayne Wolf, <i>Multiprocessor Systems-on-Chips</i> (2005) (“Wolf”)
1009	Wael Badawy, <i>System on Chip: The Challenge and Opportunities</i> , SYSTEM ON CHIP FOR REAL-TIME APPLICATIONS, 3–16 (2003) (“Badawy”)
1010	Andrei Rădulescu et al., <i>An Efficient On-Chip Network Interface Offering Guaranteed Services, Shared-Memory Abstraction, and Flexible Network Configuration</i> , 2 PROC. DESIGN, AUTOMATION AND TEST IN EUR. CONF. & EXHIBITION 878 (2004) (“Rădulescu”)
1011	A. Campbell et al., <i>Integrated Quality of Service for Multimedia Communications</i> , Computing Dept. Lancaster University (“Campbell”)
1012	D. Kandlur et al., <i>Real-Time Communication in Multihop Networks</i> , IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, VOL. 5, No. 10 (1994) (“Kandlur”)
1013	Jose Duato, et al., <i>Interconnection Networks: An Engineering Approach</i> , IEEE COMPUTER SOCIETY PRESS, Ch. 7 (1997) (“Duato”)
1014	Axel Jantsch and Hannu Tenhunen, <i>Networks on Chip</i> , Ch. 1-3 (2003) (“Jantsch”)

Exhibit No.	Description
1015	David E. Culler and Jaswinder Pal Singh, Parallel Computer Architecture, A Hardware/Software Approach, Ch. 10 (1999) (“Culler”)
1016	Hugo De Man, <i>System-on-chip design: Impact on Education and Research</i> , IEEE DESIGN & TEST OF COMPUTERS, Vol. 16, No. 3, pp. 11-19 (July-Sept. 1999) (“Man”)
1017	ETSI Technical Report, Trans European Trunked Radio (TETRA) system; Technical requirements specification, Part 2: Packet Data Optimised (PDO) systems (January 1994)
1018	U.S. Patent No. 5,828,658 (“Ottersten”)
1019	Joint Claim Construction and Prehearing Statement, <i>Network Sys. Techs., LLC v. Tex. Instruments Inc.</i> , No. 2:22-cv-00482-RWS, Dkt. 84 (E.D. Tex. Aug. 17, 2023)
1020	Pl.’s Opening Claim Construction Br., <i>Network Sys. Techs., LLC v. Tex. Instruments Inc.</i> , No. 2:22-cv-00482-RWS, Dkt. 98 (E.D. Tex. Sept. 27, 2023)
1021	E. Rijpkema <i>et al.</i> , <i>Trade-offs in the Design of a Router with Both Guaranteed and Best-Effort Services for Networks on Chip</i> , 150 IEEE PROC.-COMPUT. DIGIT. TECH. 294–302 (2003) (“Rijpkema”)
1022	M. Sgroi <i>et al.</i> , <i>Addressing the System-on-a-Chip Interconnect Woes Through Communication-Based Design</i> (2001) (“Sgroi”)
1023	Yasusi Kanada, <i>A Representation of Network Node QoS Control Policies Using Rule-based Building Blocks</i> (2000) (“Kanada”)
1024	IEEE 100, <i>The Authoritative Dictionary of IEEE Standards Terms</i> (2000) (“IEEE Dictionary”)
1025	Declaration of Gordon MacPherson regarding K. Goossens, et al., <i>Networks on Silicon: Combining Best-Effort and Guaranteed Services</i>
1026	IEEE Bibliographic Data Regarding K. Goossens, et al., <i>Networks on Silicon: Combining Best-Effort and Guaranteed Services</i>
1027	Declaration of the University of Texas at Austin Libraries regarding “Networks on Chip”
1028	SpringerLink Bibliographic Data Regarding “Networks on Chip”

Exhibit No.	Description
1029	William J. Dally and Brian Towles, <i>Route Packets, Not Wires: On-Chip Interconnection Networks</i> , Computer Systems Laboratory Stanford University (2001) (“Dally”)
1030	Defendants’ Joint Claim Construction Br. (Excerpts), <i>Network Sys. Techs., LLC v. Tex. Instruments Inc.</i> , No. 2:22-cv-00482-RWS, Dkt. 77 (E.D. Tex. Oct. 12, 2023)

I, Dr. Todd Mowry, declare as follows:

1. My name is Todd Mowry.
2. I have been retained as an expert witness on behalf of Samsung Electronics Co., Ltd., (“Samsung” or “Petitioner”) for the above-captioned Petition for *Inter Partes* Review (“Petition”) of U.S. Patent No. 7,373,449 (the “’449 patent”) (Ex. 1001). I am being compensated for my time in connection with this Petition at my standard consulting rate of \$650 per hour. My compensation is not affected by the outcome of this matter.
3. I have been asked to provide my opinions regarding whether claims 1–6 and 9–16 of the ’449 patent (the “Challenged Claims”) are invalid as obvious to a person having ordinary skill in the art at the time of the alleged invention.
4. The ’449 patent issued on May 13, 2008, from U.S. Application No. 10/530,267 (Ex-1002, “the ’267 Application”), which was filed on April 5, 2005. The ’267 Application is a national stage application (under 35 U.S.C. § 371) of PCT/IB03/04414, which was filed on October 7, 2003. The ’449 patent claims priority to a foreign patent application, EP Application No. 02079196, which was filed on October 8, 2002.
5. I am not currently, and have not at any time in the past been, an employee of Samsung. I have no financial interest in Samsung.

I. BACKGROUND AND QUALIFICATIONS

6. I am a Professor in the Department of Computer Science at Carnegie Mellon University. I also have a courtesy appointment in the Department of Electrical and Computer Engineering. I have served on the faculty of Carnegie Mellon University for twenty-six (26) years starting in 1997 through the present (2023).

7. I also served on the faculty of the University of Toronto for four (4) years between 1993 and 1997, in the Department of Electrical and Computer Engineering and a courtesy appointment in the Department of Computer Science. Prior to that appointment, I served as a Graduate Research Assistant in the Department of Electrical Engineering at Stanford University for four (4) years between 1989 and 1993. As a faculty member, I have taught and continue to teach courses and directed research in computer architecture, parallel processing, computer systems and software, compiler optimization, and operating systems.

8. I received a B.S. degree in Electrical Engineering with Highest Distinction from the University of Virginia in May 1988. I received an M.S. in Electrical Engineering from Stanford University in June 1989, and a Ph.D. in Electrical Engineering from Stanford University in March 1994.

9. I have worked in the computer industry in various capacities. I was a part-time Computer Architect and then Computer Architecture Consultant at Silicon

Graphics, Inc. in Mountain View, California (formerly MIPS Computer Systems in Sunnyvale, California) from 1989 to 1993 and 1993 to 1996, respectively. I was a Visiting Scientist at IBM in Toronto from 1996 to 2004. During that same time period (1996 to 2004), I was also a Member of the Technical Advisory Board of SandCraft, Inc. in Santa Clara, California. I was the Director of the Intel Research Pittsburgh Lab at Intel Corporation in Pittsburgh, Pennsylvania from 2004 to 2007.

10. I have authored 19 journal articles and over 60 conference papers. I am also an inventor on 5 patents.

11. I am the recipient of several honors and awards: the Arthur Samuel Thesis Award (awarded by the Stanford Computer Science department to the top two Ph.D. theses in a given year), several IBM Faculty Development Awards (1996, 1997, 1998, 2000, 2001, 2002, and 2003), several Best Paper Awards (the Second Symposium on Operating Systems Design and Implementation in 1996; the 20th International Conference on Data Engineering (ICDE) in 2004), the Alfred P. Sloan Research Fellowship (awarded to researchers in recognition of distinguished performance and a unique potential to make substantial contributions to their field), the Most Thought-Provoking Idea Award in 2004 (awarded by the Architectural Support for Programming Languages and Operating Systems (ASPLOS)), the TR100 Award in 1999 (awarded by MIT's Technology Review magazine to the top

100 most promising young innovators in science and technology), and I became an ACM Fellow in 2016.

12. I am a member of the Institute of Electrical and Electronics Engineers (IEEE) and the Association of Computing Machinery (ACM). I was the Editor-in-Chief of ACM Transactions on Computer Systems from 2013 through 2018, which is the premier journal for computer systems research. (I was an Associate Editor for the journal prior to that, since 2001.) I was the Program Chair of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) in 2010. I was the Co-Program Chair of the International Conference on Parallel Architectures and Compilation Techniques (PACT) in 2001. I have been on the programming committee in various years for ASPLOS, the International Symposium on Computer Architecture (ISCA), the International Symposium on Microarchitectures, and the Workshop on Architectural and System Support for Improving Format.

13. I have authored numerous publications relating to systems-on-chips (SOCs) comprised of multiple processors that communicate with each other via networks-on-chips (NOCs). For example, when the Association of Computing Machinery (ACM) made me an ACM Fellow in 2016, they explicitly recognized my research contributions in the area of thread-level speculation, which was a new technique specifically motivated by these architectures. As other examples, my

work in the Log-Based Architectures project proposed a novel use of SOC's where logs were captured and shipped across NOCs to other processors on the same chip, and my work on on-chip memory compression helped to reduce the load on the NOC (and also led to an interesting observation about how memory compression had an unintended consequence on NOC energy, which we were able to fix).

14. I am also the named inventor on multiple issued patents relating to processor architecture.

15. My qualifications and publications are set forth more fully in my curriculum vitae, attached as Ex. 1004.

II. MATERIALS AND OTHER INFORMATION CONSIDERED

16. In forming the opinions expressed in this Declaration, I relied upon my education and experience in the relevant field of the art and have considered the viewpoint of a person having ordinary skill in the art (POSITA) at the time of the alleged invention.

17. I have considered the materials referenced herein, including the '449 patent (Ex. 1001), the file history of the '449 patent (Ex. 1002), the Petition, and other documents listed in the Exhibit List of the Petition, including:

Description	Date of Public Availability
K. Goossens, et al., <i>Networks on Silicon: Combining Best-Effort and Guaranteed Services</i> ("Goossens2002") (Ex. 1005)	Published no later than August 7, 2002

Description	Date of Public Availability
U.S. Patent No. 5,461,611 (“Drake”) (Ex. 1006)	Filed on June 7, 1994 and issued on October 24, 1996.
K. Goossens, et al., <i>Guaranteeing the Quality of Services in Networks on Chip</i> (“Goossens2003”) (Ex. 1007)	Published no later than January 31, 2003.

18. The references listed above are prior art to the ’449 patent.

III. UNDERSTANDING OF PATENT LAW

19. I am not an attorney. For purposes of this declaration, I have been informed about certain aspects of the law that are relevant to my opinions. My understanding of the law is as listed below.

A. The Level of Ordinary Skill in the Art

20. It is my understanding that the level of ordinary skill in the art is based on a study of the patents at issue and their file histories, a study of the prior art cited therein, and knowledge of the following:

- The level of education and experience of persons actively working in the field at the time the subject matter at issue was developed;
- The types of problems encountered in the art at the time the subject matter was developed;
- The prior art patents and publications;
- The activities of others working in that field;
- Prior art solutions to those problems; and
- The sophistication of the technology at issue in this case.

It is also my understand that these factors are not exhaustive, and merely a useful guide to determine the level of ordinary skill in the art.

B. Obviousness

21. I have been informed and understand that a patent claim can be considered to have been obvious to a POSITA at the time the application was filed. This means that, even if all of the requirements of a claim are not found in a single prior art reference, the claim is not patentable if the differences between the subject matter in the prior art and the subject matter in the claim would have been obvious to a POSITA at the time the application was filed.

22. I have been informed and understand that a determination of whether a claim would have been obvious should be based upon several factors, including, among others:

- the level of ordinary skill in the art at the time the application was filed;
- the scope and content of the prior art; and
- what differences, if any, existed between the claimed invention and the prior art.

23. I have been informed and understand that the teachings of two or more references may be combined in the same way as disclosed in the claims, if such a combination would have been obvious to a POSITA. In determining whether a combination based on either a single reference or multiple references would have been obvious, it is appropriate to consider at least the following factors:

- whether the teachings of the prior art references disclose known concepts combined in familiar ways, which, when combined, would yield predictable results;

- whether a POSITA could implement a predictable variation, and would see the benefit of doing so;
- whether the claimed elements represent one of a limited number of known design choices, and would have a reasonable expectation of success by a POSITA;
- whether a POSITA would have recognized a reason to combine known elements in the manner described in the claims;
- whether there is some teaching or suggestion in the prior art to make the modification or combination of elements claimed in the patent; and
- whether the innovation applies a known technique that had been used to improve a similar device or method in a similar way.

24. I understand that a POSITA has ordinary creativity, and is not an automaton.

25. I understand that in considering obviousness, it is important not to determine obviousness using the benefit of hindsight derived from the patent being considered.

26. I understand that prior art to the '449 patent includes patents and printed publications in the relevant art that predate the Priority Date of the '449 patent.

27. I understand that certain factors—often called “secondary considerations”—may support or rebut an assertion of obviousness of a claim. I understand that such secondary considerations include, among other things, commercial success of the alleged invention, skepticism of those having ordinary skill in the art at the time of the alleged invention, unexpected results of the alleged invention, any long-felt but unsolved need in the art that was satisfied by the alleged

invention, the failure of others to make the alleged invention, praise of the alleged invention by those having ordinary skill in the art, and copying of the alleged invention by others in the field.

28. I further understand that there must be a nexus—a connection—between any such secondary considerations and the alleged invention. I also understand that contemporaneous and independent invention by others is a secondary consideration tending to show obviousness.

IV. SUMMARY OF OPINIONS

29. It is my opinion that claims 1–6, and 9–16 of the '449 patent are (i) rendered obvious by Goossens2002 in combination with Drake and (ii) rendered obvious by Goossens2003 in combination with Drake.

V. OVERVIEW OF THE TECHNOLOGY

A. Systems-on-Chip (SoC)

30. Computer systems are comprised of a number of different components, including processors (e.g, CPUs, GPUs, etc.), memory (e.g., SRAM, DRAM, etc.) and various input/output (I/O) devices. As the capacity of integrated circuit (IC) technology has improved exponentially over the years, it has enabled larger portions of computer systems to be integrated onto single chips. For example, in the early 1970s, Intel was able to integrate an entire CPU onto a single chip: e.g., the Intel 4004, which was the first commercial “microprocessor”. Compared with CPUs

during the 1960s (with their numerous separate chips on motherboards), single-chip microprocessors offered significant cost and energy savings.

31. As IC capacity has continued to increase well beyond what is needed to implement a single processor, it has enabled “Systems-on-Chip” (SoC) designs that include even more system components on a single chip: usually some combination of processors (often including multiple CPUs and possibly accelerators such as GPUs), portions of the memory hierarchy, and possibly I/O devices or controllers. SoCs were well-known by the late 1990s, and are integrated on a silicon die. *See, e.g.*, Ex-1014 (Jantsch) at vii (“During the 1990s more and more processor cores and large reusable components have been integrated on a single silicon die, which has become known under the label System on Chip (SoC).”); Ex-1008 (Wolf) at 1 (“An SoC is an integrated circuit that implements most or all of the functions of a complete electronic system.”); Ex-1009 (Badawy) at 1 (“With the increase in the number of logic gates that can be implemented on a single chip, various functionalities ... can be integrated into a single silicon chip, realizing an entire system on chip (SoC)”). Thus, SoCs can also be referenced as System on Silicon. *See, e.g.*, Ex-1009 (Badawy) at 1 (An SoC “is also known as a system on silicon”); Ex-1016 (Man) at 12 (“A SOC thus encapsulates system knowledge in a single medium—silicon—which blurs the separation between component design and system design.”).

B. Networks-on-Chip (NoC)

32. To enable the various system components within a chip for an SoC design to communicate with each other, an on-chip interconnect is needed. One such on-chip interconnect is often referred to as a “Network on Chip” (NoC). Rather than directly wiring together each pair of on-chip components via dedicated hardware FIFO buffers and wires, the idea of a NoC design is to create a more general on-chip interconnect that routes packets from sources to destinations, often traversing multiple hops through on-chip routing switches. The rationale for this NoC approach was described in the influential paper by Bill Dally and Brian Towles that appeared in the Design Automation Conference (DAC)—the premier IC technology conference—in June, 2001. Ex-1029 (Dally).

33. By integrating the interconnection network fully onto a chip, NoC designs can potentially provide lower cost, lower latency, and higher bandwidth relative to off-chip designs. To take full advantage of these potential benefits and to provide scalability as the number of on-chip components increases, NoCs have borrowed many design features from scalable interconnection networks that were designed for large-scale parallel machines in the 1990s. *See* Ex-1015 (Culler) at 749–828; Ex-1010 (Radulescu) at 277 (“Networks have been the subject of research for decades, both in the context of local and wide area networks (computer networks) [], and as an interconnect for parallel machines []. Both are very much related to on-

chip networks, and many of the results in those fields are also applicable on chip.”); Ex-1014 (Jantsch) at 86–95 (“[M]any research groups have concurrently proposed the idea of using a packet switched communication network, similar to one used in computer networks, for on chip communication.”); Ex-1013 (Duato) at vii-viii (“Originally developed for the demanding communication requirements of multicomputers, interconnection networks are beginning to replace buses as the standard system-level interconnection.”).

34. Because NoCs can have more diverse types of communicating components beyond traditional parallel machines (e.g., DMA engines transferring data between I/O and on-chip memory, video streaming to frame buffers, CPU cache misses to memory, data transfers between CPUs and GPUS, etc.), they are often extended to also support quality-of-service (QoS) across different communication streams.

C. Connections

35. In the field of networking, the term “connection” refers to a type of communication link between two devices. Although the term may mean different things across different layers of the OSI (Open Systems Interconnection) model (e.g., in the physical layer, it might refer to a wire or cable that physically connects two devices), it is commonly associated with the OSI transport layer, which focuses on end-to-end communication services. The transport layer can operate in one of

two modes: connection-oriented or connectionless. For wide area networks, an example of a connection-oriented protocol is TCP/IP (Transmission Control Protocol/Internet Protocol), and an example of a connectionless protocol is UDP (User Datagram Protocol).

36. For connection-oriented protocols such as TCP/IP, communication begins by first establishing a logical connection between the endpoints. The state that is associated with this logical connection is then used to manage the stream of communication: e.g., to provide guarantees on reliability, delivery order, quality of service, etc. Once the communication stream has completed, the connection is terminated (freeing up its state). In contrast, connectionless protocols such as UDP do not pre-establish a logical connection between endpoints: the sender simply goes ahead and sends its message to the receiver.

37. There are tradeoffs between these two approaches. Connectionless protocols generally have lower latency (since there is no need to first setup connection state before communicating), lower resource utilization (since they avoid the overhead of managing connection state), and better scalability when sending mostly small, independent messages (due to lower overhead). Connection-oriented protocols generally provide higher reliability (by using their connection state to implement error detection and correction mechanisms), and can also provide better bandwidth guarantees for multimedia communication streams (e.g., audio or video)

that have quality-of-service (QoS) requirements. Hence for wide-area networks, both approaches are supported (e.g., through TCP/IP and UDP), and programmers can choose which protocol is best suited to their application.

38. In the context of NoCs, similar principles apply regarding the differences between connection-oriented and connectionless approaches, although the hardware details and tradeoffs are different when the entire network is integrated within a single integrated circuit (IC). For example, NoCs offer the potential to dramatically reduce end-to-end communication latency and increase bandwidth relative to communicating off-chip. To keep these communication latencies low, it is important to minimize (or avoid) connection setup times. For example, for parallel applications that are running on multiple processor cores within a chip where communication occurs through a shared address space, communication consists of very short (typically 64 byte) messages that are sent between processors and various on-chip caches (a type of SRAM memory). Because these messages are latency-critical and very short, the latency overhead of setting up a connection for each message could hurt performance; hence parallel machines have traditionally used connectionless interconnects. For an SoC that includes other forms of communication (e.g., multimedia streams for audio or video), however, the quality-of-service benefits of connection-oriented communication may be quite compelling.

By the early 2000s, a number of papers had proposed the use of connections in NoCs. *See, e.g.*, Ex-1021 (Rijpkema); Ex-1022 (Sgroi).

39. A given connection can be made up of multiple channels. Connections can, for example, be “full duplex” in which there are two separate and independent channels, e.g., for transmission and reception of packets. *See, e.g.*, Ex-1017 (ETSI ETR 086-2) at 9 (“Duplex (full duplex): a mode of operation by which information can be transferred in both directions and where the two directions are independent.”); Ex-1018 (Ottersten) at 2:67-3:5 (“To provide full-duplex communication links, typically some of the communication channels are used for communication from base stations to users' remote terminals (the downlink), and others are used for communication from users' remote terminals to base stations (the uplink).”).

40. There are two different types of channels in NoCs: physical channels and virtual channels. A “physical channel” refers to a physical communication medium (e.g., sets of wires) that physically connect together the routing switches in the interconnection network. In contrast, a “virtual channel” is a logical connection that is usually time-multiplexed on top of physical channels. For example, if a physical channel between two switches has a bandwidth of 4 Gbytes/second, an NoC designer might choose to divide that evenly into two virtual channels (each with 2 Gbytes/second of bandwidth).

41. One common motivation for creating virtual channels is to avoid a problem called “fetch deadlock.” In a fetch deadlock scenario, a computer can become permanently stuck (i.e. deadlocked) because a number of outstanding requests cannot make forward progress because they are waiting for reply messages that cannot be sent due to circular dependences. The way to avoid fetch deadlock is to guarantee that reply messages can always make forward progress, and one easy way to ensure this is to have separate virtual channels in the network for requests versus replies. Therefore, even if the virtual channel for requests becomes completely congested, the separate virtual channel that is dedicated to replies can still make forward progress in delivering reply messages (even when the two virtual channels are sharing the same physical channel).

42. Virtual channels do not need to be divided evenly: for example, in the example above of a 4 Gbytes/second physical channel, the bandwidth could be divided up as 1 Gbyte/second and 3 Gbytes/second across two virtual channels. Hence the allocation of physical bandwidth resources to virtual channels is a useful mechanism for managing quality-of-service (QoS) in a NoC. By configuring the properties of virtual channels individually to match the corresponding application-level bandwidth requirements, networks can provide appropriate QoS guarantees. *See, e.g.*, Ex-1012 (Kandlur) at 1045 (“In a bidirectional connection between a pair of user entities, the message generation characteristics may differ substantially for

the two directions. Hence, it is preferable to restrict the real-time channel to unidirectional communication, and a bidirectional connection can be composed from a pair of channels.”); Ex-1011 (Campbell) at 10 (“For example, in a videophone connection consisting of separate video and audio channels, a slightly lower commitment probability for the video channel would be appropriate.”).

D. Admission Control

43. To avoid oversubscribing its resources, networks often use some form of admission control to limit or manage the entry of new messages into the network. As a rough analogy, if we think of messages on networks as being similar to cars on highways, the idea is to avoid having so much congestion that traffic slows considerably (or perhaps even grinds to a halt due to gridlock). To avoid this problem on highways, some cities use traffic lights on highway on-ramps to limit the amount of traffic during rush hour (to help keep highway traffic flowing). Networks use a similar approach with admission control mechanisms.

44. Admission control is especially important if a network supports quality-of-service guarantees. For example, if an SoC is already streaming a video chat across its NoC that requires more than half of its available bandwidth, it may not be possible to launch a second video stream without severely compromising the performance of one or both of the video streams. Hence the network might prefer to reject or delay the start of a second video stream until enough resources are freed

up so that it can meet the quality-of-service requirements. The network can, for example, monitor resources to perform admission control on individual connections and channels based on quality-of-service requirements. *See, e.g.,* Ex-1011 (Campbell) at 11 (“Each layer in the system contains an *admission control* module which determines whether or not the request should be considered.”).

VI. THE '449 PATENT

45. The '449 patent issued on May 13, 2008, from U.S. Application No. 10/530,267 (Ex-1002, “the '267 Application”), which was filed on April 5, 2005. The '267 Application is a national stage application (under 35 U.S.C. § 371) of PCT/IB03/04414, which was filed on October 7, 2003. The '449 patent claims priority to a foreign patent application, EP Application No. 02079196 (the “EP Application”), which was filed on October 8, 2002, but the EP Application does not disclose all of the limitations of the claims.

46. I understand that, for the '449 patent to be entitled to the earlier priority date, the EP Application must disclose all of the '449 claim elements. I have reviewed the EP Application, and in my opinion, the Challenged Claims do not have support in the EP Application. Specifically, in my opinion, the EP Application does not disclose the claimed “communication manager” from Claims 2, 6, and 10, or the functions performed by the “communication manger,” such as “forward[ing] said request (REQ) for a connection with communication channels each having a specific

set of connection properties to said resource manager (RM)” (claim 6) and “forwarding the request to a resource manager” (claim 10 and claims that depend therefrom). *See* EX-1002 (’449 FH) at 121–169. Additionally, the EP Application fails to disclose the claimed “resource manager” from claims 1, 10, and the dependent claims of the ’449 patent. The EP Application also does not disclose the claimed functions of the “resource manager,” such as “determin[ing] whether the at least one connection with said communication channels is available based on the connection properties being available” (claim 1), “determining whether a target connection with the desired connection properties is available” (claim 10), and “responding with the availability of the requested connection to said communication manager” (claims 6 and 10), among others. Ex-1002 (’449 FH) at 121–169.

A. Claims

47. The ’449 patent has 18 claims, including two independent claims numbered 1 and 10. Ex. 1001, 18:32-20:38. Claims 1–6 and 9–16 are the Challenged Claims. Exemplary Claims 1 and 10 are reproduced below:

1. Integrated circuit comprising a plurality of processing modules (M; I; S; T), network (N; RN), and at least one resource manager (RM) for managing resources of the network (N), wherein the modules, network and resource manager are disposed on a chip and arranged for providing at least one connection between a first and at least one second module, wherein the at least one connection comprises a set of communication channels each having a set of connection properties, the connection properties of the different communication channels of said connection being adjustable independently, wherein said connection

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, wherein the resource manager determines whether the at least one connection with said communication channels is available based on the connection properties being available.

Ex. 1001 ('449 patent), 18:32-48.

10. 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:

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;

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 ('449 patent), 19:21-20:6.


B. Summary of the Specification

48. As the specification explains, the '449 patent “relates to an integrated circuit having a plurality of processing modules and a network arranged for providing connections between processing modules.” Ex-1001 ('449 patent) at 1:5–11. Specifically, the '449 patent explains that “[i]t is an object of the invention to provide ... more effective usage of the properties of the network.” Ex-1001 ('449 patent) at 7:6–11.

49. The '449 patent describes that in the “transport layer” of the network, the processing modules communicate over “connections,” which are described as “a set of channels, each having a set of connection properties, between a first module and at least one second module.” Ex-1001 ('449 patent) at 8:48–52. The '449 patent further describes that the set of channels may comprise a “request channel” and a “response channel.” Ex-1001 ('449 patent) at 8:52–60. “To allow more flexibility in configuring connections, and, hence, better resource allocation per connection,” the '449 patent describes that “the outgoing and return parts of connections can be configured independently,” such that they have different connection properties. Ex-1001 ('449 patent) at 11:4–55. The “connection properties” provided by these channels may include, for example, throughput and latency. Ex-1001 ('449 patent) at 8:62–9:3.

50. The '449 patent also describes that “connections ... must be first created or established with the desired properties before being used.” Ex-1001 ('449 patent) at 10:62–64. To establish such connections, the '449 patent describes that a “communication manager (CM) request[s] a connection with a set of properties between two modules from the resource manager RM.” Ex-1001 ('449 patent) at 10:19–24. The '449 patent explains that, because “[t]he properties require resources to be implemented[,] ... a connection requiring some properties is opened or not depending on the availability of these resources.” Ex-1001 ('449 patent) at 10:27–31. As such, the the resource manager RM “enquires if such a connection based on these resources is possible.” Ex-1001 ('449 patent) at 10:24–26. If the resources are available, the resource manager RM “establishes a connection.” Ex-1001 ('449 patent) at 10:41–57.

51. Figure 3, shown below, illustrates an exemplary embodiment of the system of the '449 patent.



Ex-1001 ('449 patent) at Fig. 3.

C. Summary of the Prosecution History

52. I have reviewed the file history of the '449 patent.

53. The '267 application was originally filed with 7 claims, of which two were independent. The originally filed independent claims—which correspond to issued claims 1 and 10—contained a subset of the limitations found in the issued claims. Application claim 1 is representative:

1. Integrated circuit comprising a plurality of processing modules (M; I; S; T) and a network (N; RN) arranged for providing at least one connection between a first and at least one second module,

wherein the at least one connection comprises a set of communication channels each having a set of connection properties, the connection properties of the different communication channels of said connection being adjustable independently,

wherein said connection supports transactions comprising outgoing messages from the first module to the second module and/or return messages from the second module to the first module.

Ex-1002 ('449 FH) at 105.

54. On March 2, 2007, the examiner rejected all pending claims as obvious over U.S. Patent No. 4,807,118 ("Lin") in view of U.S. Patent App. Pub. No. 2003/0182419 ("Barr"). Ex-1002 ('449 FH) at 53–59. The examiner found that the first and third limitations were disclosed by Lin and that the second limitation, a connection comprising a set of channels with independently adjustable connection properties, was disclosed by Barr. Ex-1002 ('449 FH) at 53–59.

55. On May 25, 2007, the applicants submitted an amendment with minor changes to the independent claims and four new dependent claims. Ex-1002 ('449 FH) at 43–47. In their remarks, the applicants traversed the examiner's rejection on the basis that neither Lin nor Barr was directed to “an Integrated circuit comprising a plurality of processing modules (M; I; S; T) and a network (N; RN).” Ex-1002 ('449 FH) at 48. The applicants further argued that the combination did not disclose “the processing modules and the network being disposed on the chip,” as recited in newly added dependent claims 9 and 11. Ex-1002 ('449 FH) at 48.

56. On August 23, 2007, the examiner issued another office action rejecting all of the pending claims as unpatentable over U.S. Patent App. Pub. No. 2004/0103230 (“Emerson”) in view of U.S. Patent No. 6,769,046 (“Adams”). Ex-1002 ('449 FH) at 31–37. The examiner found that Emerson discloses a “system on chip (i.e. SoC) comprising numbers of processing modules communicating within the chip ... the message between modules being exchanged over a connection via a network.” Ex-1002 ('449 FH) at 33. In addition, the examiner found that Adams discloses connections comprising a set of channels with independently adjustable connection properties. Ex-1002 ('449 FH) at 33.

57. On November 20, 2007, the applicants submitted an amendment incorporating into independent application claims 1 and 7 limitations from dependent application claims 2 and 3 pertaining to a “resource manager” and adding

seven new dependent claims. Ex-1002 ('449 FH) at 20–24. In their remarks, the applicants traversed the examiner's rejection, arguing that Emerson's arbiter, which the examiner had mapped to the claimed resource manager, "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." Ex-1002 ('449 FH) at 25–26 (emphasis in original).

58. The application was allowed on January 9, 2008, with the examiner noting that the cited prior art did not teach that "the resource manager determines whether the at least one connection with the communication channels is available based on the connection properties being available." Ex-1002 ('449 FH) at 13–15.

VII. LEVEL OF ORDINARY SKILL IN THE ART

59. I understand that there are multiple factors relevant to determining the level of ordinary skill in the pertinent art, including the educational level of active workers in the field at the time of the alleged invention, the sophistication of the

technology, the type of problems encountered in the art, and the prior art solutions to those problems.

60. In determining the characteristics of a hypothetical person of ordinary skill in the art of the '449 patent at the time of the alleged invention, I considered several things, including the type of problems encountered in this field, and the rapidity with which innovations were made. I also considered the sophistication of the technology involved, and the educational background and experience of those actively working in the field, and the level of education that would be necessary to understand the '449 patent. Finally, based on my own experience, I placed myself back in the relevant periods of time, and considered the state of the art and the level of skill of the engineers working in this field in those times.

61. It is my opinion that a POSITA at the time of the '449 patent 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. Additional education might compensate for less experience, and vice-versa.

62. I understand that Patent Owner has proposed a somewhat different level of skill in the art in parallel district court proceedings:

NST's expert, Dr. Chmelar, opines that a person of ordinary skill in the art ("POSITA") at the time of the invention of the Asserted Patents is an individual having at least: a Bachelor of Science (or equivalent) degree in electrical engineering, computer

engineering, or a related field, and 2-3 years of work experience in very large-scale integrated (“VLSI”) systems, such as application-specific integrated circuits (“ASICs”), application-specific standard parts (“ASSPs”), system-on-chip (“SoC”), or field-programmable gate arrays (“FPGAs”), and has an understanding of on-chip interconnection networks.

Ex-1020 at 7. My opinions provided in this Declaration would not change under the level of ordinary skill in the art proposed by Patent Owner.

VIII. OVERVIEW OF THE PRIOR ART

A. Overview of Goossens2002

63. Goossens2002 was authored by a group from Philips Research Laboratories that included one of the named inventors of the ’449 patent, Kees Goossens. Ex-1005 (Goossens2002) at 1. Goossens2002 is directed to “systems on silicon,” also known as a systems on chip, or “SoCs.” Ex-1005 (Goossens2002) at 1. Goossens2002 describes that such systems on silicon are potentially comprised of “hundreds of IPs,” or intellectual property blocks, and is concerned with the communication among those IPs. Ex-1005 (Goossens2002) at 1. Goossens2002 “propose[s] a *network on silicon* (NOS) to implement the communication between these IPs.” Ex-1005 (Goossens2002) at 1.

64. The network on silicon proposed by Goossens2002 is a packet-switched router network that implements “guaranteed services.” Ex-1005 (Goossens2002) at 1. Goossens2002 describes that “some IPs have inherent performance requirements, such as a minimum throughput (for real-time streaming

data), or bounded latency (for interrupts)” and that “[a] NOS with guaranteed services eases their integration” into larger systems on silicon. Ex-1005 (Goossens2002) at 2. The guaranteed services in Goossens2002’ system are provided over “*connections*,” which are combinations of multiple channels between IPs that can be set up and removed by the IPs. Ex-1005 (Goossens2002) at 2–3. Goossens2002 describes that the component channels “can have a combined throughput and latency guarantee (GT), or be best effort (BE), with an unspecified finite latency upper bound.” Ex-1005 (Goossens2002) at 3.


65. Goossens2002 further notes that “[a] service request may be granted or denied by the network” and suggests that this functionality is implemented in a “quality-of-service manager” that “reconfigures” the system on silicon and “steers” the network on silicon. Ex-1005 (Goossens2002) at 3.

B. Overview of Drake

66. Drake is titled “Quality of Service Management for Source Routing Multimedia Packet Networks,” and was originally assigned to International Business Machines Corporation. Ex-1006 (Drake) at Cover. Drake describes that multimedia information—“particularly audio and full motion video”—that is communicated over a network “has unique properties that impose stringent requirements on the communications network.” Ex-1006 (Drake) at 1:22–24. Drake explains that these requirements include the availability of “high bandwidth” (*i.e.*, throughput) and that

“the data is delivered with (a priori) bounded delay” (*i.e.*, latency). Ex-1006 (Drake) at 1:34–46. Drake further describes the need for a system and method to ensure that these requirements can be satisfied in packet-switched networks. Ex-1006 (Drake) at 1:52–57, 2:14–21.

67. To address this need, Drake describes a quality of service management system that receives requests for “QoS connections” with specified properties, determines whether the required network resources are available, and reserves those resources for the requesting station. Ex-1006 (Drake) at Abstract. An exemplary embodiment of Drake’s system is depicted in Figure 1 below.



Ex-1006 (Drake) at Fig. 1. The components of Drake's system include a "QoS requestor manager," which "assemble[s] a request to reserve a QoS connection," and a "QoS allocator," which receives the request, "determine[s] if a path exists ... which satisfies the requested level of QoS," and assigns the available resources to the requestor. Ex-1006 (Drake) at 4:26–56.

68. Drake further describes that the component channels of the QoS connection—the "Forward Direction" channel from the source to the target, and the "Backward Direction" channel from the target to the source—have separately and independently configured properties. Ex-1006 (Drake) at 5:30–47.

C. Overview of Goossens2003

69. Goossens2003 was authored by a group from Philips Research Laboratories that included both of the named inventors of the '449 patent, Kees Goossens and Andrei Rădulescu. Ex-1007 (Goossens2003) at 61. Goossens2003 is directed to a "network on chip (NOC)"—the *ÆTHEREAL* NOC—for use in providing communications between "intellectual property blocks (IP)" in a system-on-chip. Ex-1007 (Goossens2003) at Abstract, 63, 74. Goossens2003 explains that such a NOC "integrates heterogeneous IPs in a standard fashion." Ex-1007 (Goossens2003) at 66.

70. Goossens2003 describes that "[t]he *ÆTHEREAL* NOC offers differentiated services with *connections*" and explains that a "connection describes

communication between one master and one or more slaves, with an associated service level.” Ex-1007 (Goossens2003) at 71. The services provided by these connections, which Goossens also refers to a “connection properties,” include “data integrity,” “transaction ordering,” “transaction completion,” “connection flow controller,” and “connection throughput, latency, and jitter.” Ex-1007 (Goossens2003) at 72–74. Goossens2003 further describes that “[c]onnections must be created stating the requested service level” and that “[t]he NOC either accepts or rejects the request for the connection.” Ex-1007 (Goossens2003) at 71. This is described as the “negotiation phase” of a “resource management” process. Ex-1007 (Goossens2003) at 71; *see also* Ex-1007 (Goossens2003) at 65 (“Resource management depends on two phases: negotiation to obtain resources, followed by a steady state in which allocated resources are used.”). Goossens2003 also describes that the “request and response” channels of a connection “can be independently configured ... allowing more fine-grained resource management.” Ex-1007 (Goossens2003) at 74.

IX. SPECIFIC GROUNDS FOR PETITION

A. Ground I: Claims 1–6 and 9–16 Are Rendered Obvious by Goossens2002 in View of Drake

71. In my opinion, claims 1–6 and 9–6 are rendered obvious by the combination of Goossens2002 and Drake.

72. As a threshold matter, a POSITA would have been motivated to combine the disclosures of these two references because both relate to the provision of guaranteed services in a communication network, and Drake provides additional implementation details that would aid a POSITA in creating an SoC with improved performance and efficiency.

73. As I summarized above, Goossens2002 describes a network on chip architecture for SoCs. Ex-1005 (Goossens2002) at 1. One of the primary focuses of the paper is the provision of guaranteed services, which Goossens2002 describes as “essential” for supporting the communication requirements of certain types of IP blocks that might be incorporated into an SoC. Ex-1005 (Goossens2002) at 1. The paper gives as examples IP blocks involved in real-time streaming, which require a guarantee that there will be sufficient bandwidth available (“minimum throughput”), and IP blocks that issue interrupts, which require a guarantee that the interrupts will be delivered within a specified amount of time (“bounded latency”), and multimedia transmission, which requires the synchronization of audio and video streams. Ex-1005 (Goossens2002) at 1–2. In addition to minimum throughput and bounded latency, Goossens2002 describes guaranteed services that include “uncorrupted transmission, lossless transmission, in-order delivery,” and “jitter (delay variation).” Ex-1005 (Goossens2002) at 2.

74. Goossens2002 describes that these guaranteed services are provided through “connections” that can be established by the IP blocks when needed, and removed when they are no longer necessary. Ex-1005 (Goossens2002) at 2–3. As I explained in the technology background section above, by the early 2000’s, connections were a well-known mechanism for providing guaranteed services.

75. However, Goossens2002 does not include certain details regarding the implementation of its guaranteed services and connections. For example, the paper describes that IP blocks can make requests for guaranteed services (and the establishment of corresponding connections) but does not describe how those requests are formulated or where they are sent. Goossens2002 also states that these requests may be either granted or denied but does not describe how the system makes this determination. Ex-1005 (Goossens2002) at 3. The paper does refer to a “quality-of-service manager,” which a POSITA would have recognized as a component involved in this process, but it does not provide any implementation details regarding this component or its process for determining whether to accept or reject requests for guaranteed services. Ex-1005 (Goossens2002) at 3.

76. Drake provides details regarding the implementation of these features and a POSITA would have found Drake’s known solution for “ensur[ing] selected levels of quality of services” helpful in implementing an SoC like the one described by Goossens2002. See Ex-1006 (Drake) at Abstract. A POSITA would have been

particularly interested in Drake's disclosure of a "QoS requestor manager" to formulate requests for guaranteed services with particular properties and a "QoS allocator" to determine whether the network resources required for the requested services are available. *See* Ex-1006 (Drake) at 2:14–40, 4:17–56.

77. Like Goossens2002, Drake describes providing guaranteed services for certain kinds of network communications, including guarantees related to throughput and latency. Drake describes that these services are particularly advantageous for multimedia applications, which require large amounts of bandwidth to transmit video data and bounded delay in the delivery of the data. Ex-1006 (Drake) at 1:22–46. Also, like Goossens2002, Drake describes providing these guaranteed services using "connections." *See, e.g.,* Ex-1006 (Drake) at 4:26–29.

78. Drake recognizes that a critical element in providing such guaranteed services is ensuring the availability of the required network resources. To do so, Drake describes a quality of service management system that includes a "QoS requestor manager" to formulate requests for guaranteed services with particular properties and a "QoS allocator" to determine whether the network resources required for the requested services are available. Ex-1006 (Drake) at Abstract, 2:14–40, 4:17–56, Figs. 1, 4–5. When the required resources are available, they are reserved for the pending request in order to ensure their availability, and when they are not available the request is denied. Ex-1006 (Drake) at Abstract. Drake explains

that “[t]he intent of such a technique is to deny network access to those streams whose level of quality of service requirements are not currently available.” Ex-1006 (Drake) at 2:18–21.

79. In my opinion, a POSITA would have understood that using Drake’s quality of service management process would improve the performance of Goossens2002’s network on silicon by ensuring that the network resources required to provide guaranteed services are actually available, and that the network will be able to fulfill its commitments to the IPs that have been granted such services. In addition, in my opinion, a POSITA would have understood that using Drake’s quality of service management process would improve the efficiency of Goossens2002’s network on silicon by allowing IP blocks to reserve only those resources that they actually need. For example, Drake describes that different QoS parameters and can be specified—and therefore different resources reserved—for an output channel and an input channel. Ex-1006 (Drake) at 5:30–50. This allows the reservation of resources to be tailored to the needs of a given request, increasing the pool of resources available for other IP blocks, and minimizing waste. To use examples provided by Drake, the playback of stored video only requires high throughput in a single direction, while two-way video conferencing requires high throughput in both directions. Ex-1006 at 1:34–44.

80. In my opinion, a POSITA would also have reasonably expected a combination of the disclosures of Goossens2002 and Drake to be successful. For instance, a POSITA would have recognized that these two references are both in the same field of study—guaranteed services for data communications in packet-based networks—and both are directed to ensuring that sufficient resources are available for the communication of multimedia content, such as audio and video. *See, e.g.*, Ex-1005 (Goossens2002) at 2 (describing “audio and video streams”); Ex-1006 (Drake) at 1:22–24 (describing “audio and full motion video”). A POSITA would also have recognized that such a combination would have involved nothing more than combining prior art elements according to known methods to yield predictable results. Drake provides additional details about a system for providing guaranteed services that is consistent with and would have worked with the system of Goossens2002. As such, a POSITA would have reasonably expected that use of Drake’s known solution to ensuring the availability of network resources for guaranteed services would improve Goossens2002’s network on silicon in the same way.

81. In addition, while Drake describes its QoS management system in the context of local area networks, a POSITA would have recognized that the advantages and benefits of Drake’s system are applicable to a network on silicon like the one described by Goossens2002. For instance, Goossens2002 itself notes

that guaranteeing services in a network on silicon requires careful management of resources. See Ex-1005 (Goossens2002) at 3 (“a NOS that offers guaranteed services must accurately model its resources to offer those services). Drake’s QoS management process serves this very purpose and its more efficient utilization of network resources is well-suited for such a resource-constrained environment.

82. Furthermore, the parallels between mature off-chip networks and newer on-chip networks would have been well-known to POSITA. As I explained above, NoCs have borrowed many design features from scalable interconnection networks that were designed for large-scale parallel machines in the 1990s. For example, as described in one text devoted to networks-on-chip “[d]esigning a network for on-chip communication is similar to the problem of designing a local area computer network with some computing and communication requirements.” Ex-1014 (Jantsch) at 102. Another chapter in the same text explained that addressing issues related to network resource allocation—as Goossens2002 and Drake do—is substantially similar in on-chip and off-chip networks: “The resource allocation, optimization of network usage and verification of performance and correctness are the main problems that are basically similar to what distributed and parallel system designers have to face.” Ex-1014 (Jantsch) at 33. Indeed, the ’449 patent itself acknowledges that many of the issues faced by networks on chip had been studied

in the context of off-chip networks and that results of those studies are applicable to on-chip networks as well. Ex-1001 ('449 patent) at 1:50–66.

1. Claim 1

83. In my opinion, Goossens2002 in view of Drake renders claim 1 obvious.

a. Claim 1[pre]: “Integrated circuit comprising”

84. I was not asked to analyze whether the preamble of Claim 1 is limiting. In any event, in my opinion, Goossens2002 discloses an integrated circuit. Goossens2002 is directed to communications in “systems on silicon (SOSs),” also known as systems on chip (SoCs). Ex-1005 (Goossens2002) at 1 (“In this section we consider how a SOS containing hundreds of IPs can be designed....”); Ex-1009 (Badawy) at 1 (An SoC “is also known as a system on silicon”).

85. As I described above in the background of the technology (*see* above Section V.A) a POSITA would have understood that a SOS or SOC is implemented on a single integrated circuit chip, and thus all of the components of the system described by Goossens2002 are on the same chip. *See, e.g.*, Ex-1014 (Jantsch) at vii (“During the 1990s more and more processor cores and large reusable components have been integrated on a single silicon die, which has become known under the label System on Chip (SoC).”); Ex-1016 (Man) at 12 (“A SOC thus encapsulates system knowledge in a single medium—silicon—which blurs the separation

between component design and system design.”); Ex-1008 (Wolf) at 1 (“An SoC is an integrated circuit that implements most or all of the functions of a complete electronic system.”); Ex-1009 (Badawy) at 1 (“With the increase in the number of logic gates that can be implemented on a single chip, various functionalities ... can be integrated into a single silicon chip, realizing an entire system on chip (SoC)”). Indeed, Goossens2002 refers to “the implementation of the physical *on-chip* wires.” Ex-1005 (Goossens2002) at 2.

86. Additionally, as I discuss below with respect to claim elements 1[a]–[b], Goossens2002’s system on silicon includes IP blocks interconnected by a network on silicon. *See* Below Claim Elements 1[a]–1[b]. Therefore, in my opinion, Goossens2002 discloses an integrated circuit, including under the construction I understand that Patent Owner (PO) has proposed in parallel district court proceedings: “interconnected circuitry on a chip.” *See* Ex-1019 at 9.

b. Claim 1[a]: “a plurality of processing modules (M; I; S; T)”

87. In my opinion, Goossens2002 discloses this limitation. For example, Goossens2002 describes that systems on silicon include multiple intellectual property blocks or “IPs.” Ex-1005 (Goossens2002) at 1 (“a SOS contains hundreds of IPs that are bought or premade”). Goossens2002 also describes that these IPs may be processors that perform computations. Ex-1005 (Goossens2002) at 1 (referring to “computation (IPs)”), 2 (“computation is mapped to functional IPs (e.g.

tasks to processors)”). Indeed, in my opinion, this discussion is consistent with the specification of the ’449 patent, which describes that the “modules” of its disclosed embodiments “can be so-called intellectual property blocks IPs (computation elements, memories or a subsystem which may internally contain interconnect modules).” Ex-1001 (’449 patent) at 9:42–46. Therefore, in my opinion, a POSITA would have understood that the IPs in Goossens2002’s system on silicon are processing modules.

c. Claim 1[b]: “network (N; RN)”

88. In my opinion, Goossens2002 discloses this limitation. For example, Goossens2002 describes “a network on silicon (NOS).” Ex-1005 (Goossens2002) at 1 (“We propose a network on silicon (NOS) to implement the communication between these IPs.”); *see also* Ex-1005 (Goossens2002) at 1 (“We advocate a network on silicon (NOS) as a hardware architecture to implement communication between IP cores ... and as a software model in the form of a protocol stack to structure the programming of NOSs”). Goossens2002 further describes that this NOS is comprised of a “router network based on packet switching.” Ex-1005 (Goossens2002) at 1; *see also* Ex-1005 (Goossens2002) at 1 (describing “network topology and router layout”), 2 (describing “a NOS based on packet-switching,” “NOS infrastructure (such as routers and protocol stack),” 3 (describing a “NOS hardware architecture based on a packet-switched router network”). Goossens2002

also explains that the NOS implements the various layers of the “OSI [Open Systems Interconnection] protocol stack,” including the “physical layer, “data-link layer,” “network layer,” “transport layer,” “session layer,” presentation layer,” and “application layer. Ex-1005 (Goossens2002) at 2.

89. I understand that Patent Owner has advocated in parallel district court proceedings for a construction of “network” as “a multi-hop interconnection between modules separated by one or more network nodes.” Ex-1019 at 9. To the extent Patent Owner urges such a construction here, in my opinion, a POSITA would have understood that Goossens2002 discloses such a network. Specifically, Goossens2002 explains that a NOS packet traverses a number of routers, and refers to each such traversal as a “hop.” Ex-1005 (Goossens2002) at 2. As was well-understood to a POSITA, Goossens2002’s routers enable the packet communication from source to destination, acting as network nodes. *See, e.g.*, Ex-1023 (Kanada) at 161 (describing “network nodes, e.g., QoS-ready routers”). Goossens2002 explains that “[i]n a NOS a packet must traverse a number of routers.” Ex-1005 (Goossens2002) at 2. Goossens2002 also explains that there are delays associated with each *hop*, for example, with static scheduling. *See* Ex-1005 (Goossens2002) at 2 (“When scheduling statically, a packet need not incur more than a pipeline delay per hop.”). Therefore, in my opinion, a POSITA would have understood that Goossens2002’s routers are “network nodes,” and because the packets must traverse

a number of routers (*i.e.*, multiple hops) means that Goossens2002's NOS is a "multi-hop interconnection between modules separated by one or more network nodes."

d. Claim 1[c]: "at least one resource manager (RM) for managing resources of the network (N)"


90. In my opinion, this limitation is disclosed by Goossens2002 and/or rendered obvious by Goossens2002 in view of Drake. For example, Goossens2002 describes that its system on silicon includes a "quality-of-service manager" (*i.e.*, resource manager) that can "observe and steer" the NOS. Ex-1005 (Goossens2002) at 3. As a component of Goossens2002's system on silicon, which is an integrated circuit, a POSITA would have recognized that the quality-of-service manager is a circuit. In my opinion, a POSITA would also have recognized that this functionality involves managing resources of the network. Indeed, Goossens2002 explains that "[a]t run-time an application or quality-of-service manager *reconfigures the SOS*," and that these QoS managers "effectively observe and steer [the] NOS." Ex-1005 (Goossens2002) at 3. Goossens2002 explains that it can offer guaranteed services, *e.g.*, based on QoS, by "accurately model[ing] its resources to offer those services"—*i.e.*, through use of the QoS manager. Ex-1005 (Goossens2002) at 3. I understand that Patent Owner has proposed that "resource manager" means "circuit that manages network resources." EX-1019, at 13. In my opinion, Goossens2002 discloses such a resource manager.

91. Additionally, Drake similarly discloses a “resource manager” in the form of a “quality of service (QoS) allocator.” Ex-1006 (Drake) at 4:22–25; *see also* Ex-1006 (Drake) at 4:66–5:14, Fig. 1. Drake further describes the QoS allocator as managing resources of the network. Specifically, the QoS allocator receives requests for connections with certain properties and examines the resources of the network to determine whether those requests can be satisfied. Ex-1006 (Drake) at 4:49–54 (“The overall function of allocator 20 is to receive the QoS request from station 10, examine the resources of LAN 17, and determine if a path exists in LAN 17 between station 10 and station 29 which satisfies the requested level of QoS.”); *see also* Ex-1006 (Drake) at Abstract (“An allocator in the management system determines if the requested facilities are available and, if so, reserves these facilities for the requesting station.”), 2:14–21, 2:65–3:6, 10:41–51, 10:52–11:48, Figs. 4–5. Drake also describes those connection properties as “quality-of-service parameters that describe the service expected from the network,” which include throughput and end-to-end delay (*i.e.*, latency). Ex-1006 (Drake) at 5:30–50. In order to enable this functionality, Drake describes that the QoS allocator includes a “QoS allocation database” that stores information regarding the resources of the network as well as information regarding the existing connections using those resources. *See* Ex-1006 (Drake) at 5:8–14; *see also* Ex-1006 (Drake) at 2:32–41, 5:14–23, Figs. 1–2. In my

opinion, Drake discloses a resource manager under Patent Owner's proposed construction.

92. Indeed, Drake's QoS allocator provides the same functionality as the "resource manager" that is described in the specification of the '449 patent. Specifically, like Drake, the '449 patent describes that its resource manager receives requests for connections with certain properties, "enquires if such a connection based on these resources is possible," and "reserves the available connection i.e. the required resources." Ex-1001 ('449 patent) at 10:19–57. Additionally, the "quality-of-service parameters" described by Drake—specifically throughput and end-to-end delay (i.e., latency)—are consistent with the "connection properties" identified by the '449 patent. See, e.g., Ex. 1001 ('449 patent) at 8:62–9:3 (stating "connection properties may include ... throughput ... [and] latency"), 9:52–55 ("Connections are introduced to describe and identify communication with different properties, such as guaranteed throughput, bounded latency...."), 13:5–9 (similar).

93. Drake's Figure 1 annotated below shows the QoS allocator:



Ex-1006 at Fig. 1 (annotated).

94. For the reasons I discussed above, in my opinion, a POSITA would have found it obvious to look to Drake's disclosures regarding a QoS allocator for details regarding the implementation of Goossens2002's quality-of-service manager, and would have had a reasonable expectation of success in doing so. See above Section IX.A.

- e. **Claim 1[d]: “wherein the modules, network and resource manager are disposed on a chip and arranged for providing at least one connection between a first and at least one second module”**

95. In my opinion, this limitation is disclosed by Goossens2002, and additionally rendered obvious by Goossens2002 in view of Drake. As I discussed above with respect to claim elements 1[pre]–1[c], Goossens2002 describes a system on silicon that includes multiple Ips (*i.e.*, modules), a NOS (*i.e.*, network), and a “quality-of-service manager” (*i.e.*, resource manager), and in my opinion, a POSITA would have understood that all of these components are on the same chip. *See* above Elements 1[pre]-1[c]. Indeed, as I previously discussed in the background of the technology (*see* above Section V.A), the hallmark of a system on silicon like the one described by Goossens2002 is that the system and all of its components are self-contained within the same chip.

96. Additionally, as I discussed above, Drake discloses a QoS allocator (*i.e.*, resource manager), and in my opinion, a POSITA would have found it obvious to look to Drake’s disclosures regarding a QoS allocator and implement such functionality in Goossens2002’s quality-of-service manager.

97. Goossens2002 also discloses that its network on silicon is arranged to provide “connections” between the various IPs in its system on silicon. For example, Goossens2002 describes that “[t]he aim of the *transport layer* is to provide reliable end-to-end services over *connections*.” Ex-1005 (Goossens2002) at 2; *see also* Ex-

1005 (Goossens2002) at 3 (describing that “Ips can set up and remove connections”). The Ips on either end of these connections are the first and second modules. Drake similarly describes “connections” between source and target stations (*i.e.*, first and second processing modules), which it also refers to as “QoS connections,” “QoS streams” and “QoS data streams.” *See, e.g.*, Ex-1006 (Drake) at 4:26–29 (“[S]tation 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.”), 7:45–8:5 (“QoSStream 42–Represents a QoS reserved end-to-end data stream”), 14:2–7 (“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**.”), 15:28–31 (“[T]he garbage collection process of FIG. 11 serves to de-allocate resources of those **connections** for which a refresh announcement message is not received at regular intervals.”); *see also, e.g.*, Ex-1006 (Drake) at Abstract, 4:54–55, 11:17–21, 11:49–55, Fig. 2.

98. Therefore, in my opinion, Goossens2002 alone or in combination with Drake, discloses this limitation.

- f. Claim 1[e]: “wherein the at least one connection comprises a set of communication channels each having a set of connection properties, the connection properties of the different communication channels of said connection being adjustable independently”**

99. In my opinion, this limitation is disclosed by Goossens2002, and additionally rendered obvious by Goossens2002 in view of Drake.

100. Connections Having a Set of Communication Channels.

Goossens2002 describes that its connections comprise a set of communication channels. Specifically, Goossens2002 explicitly states that a “connection combines multiple channels.” Ex-1005 (Goossens2002) at 3. Goossens2002 also explains that these connections enable two-way communications between the processing modules over one-way channels. Ex-1005 (Goossens2002) at 2–3 (describing “simplex” channels and “full duplex connections”). A POSITA would understand that Goossens2002’s discussion of “full duplex” connections refers to two separate and independent channels, e.g., for simultaneous transmission and reception of packets. *See, e.g.*, Ex-1024 (IEEE Dictionary) at 344 (“duplex (1) A simultaneous, two-way, independent transmission in both directions. Synonym: full duplex.”), 1053 (“simplex (1) A communication channel that permits information transfer in one direction only. Duplex channels consist of two simplex channels simultaneously operating in opposing directions.”); Ex-1017 (ETSI ETR 086-2) at 9 (“Duplex (full duplex): a mode of operation by which information can be transferred in both

directions and where the two directions are independent.”); Ex-1018 (Ottersten) at 2:67-3:5 (“To provide full-duplex communication links, typically some of the communication channels are used for communication from base stations to users' remote terminals (the downlink), and others are used for communication from users' remote terminals to base stations (the uplink).”). Therefore, in my opinion, a POSITA would have understood that Goossens2002's connections comprise “a set of communication channels,” specifically one that provides a communication pathway in a forward direction (*e.g.*, an output channel) and another that provides a communication pathway in a reverse direction (*e.g.*, an input channel). Indeed, as I discussed above in the background of the technology, full duplex operations involve links such that information can be transferred in both directions. *See* above Section V.C (Connections).

101. **Channels Having Connection Properties.** Goossens2002 describes that each of its channels has certain connection properties. For example, Goossens2002 explains that channels can have combined throughput and latency guarantees or instead provide best effort transmission and delivery. *See* Ex-1005 (Goossens2002) 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.”); *see also* Ex-1005 (Goossens2002) at 2–3 (describing properties including “uncorrupted transmission, lossless transmission, in-order delivery, throughput,

jitter (delay variation), and latency” and “flow control and reordering for lossless in-order (multicast) transmission”). Notably, these types of connection properties are the same properties described in the specification of the ’449 patent. *See* Ex-1001 (’449 patent) at 9:52–55 (“Connections are introduced to describe and identify communication with different properties, such as guaranteed throughput, bounded latency and jitter, ordered delivery, or flow control.”); *see also* Ex-1001 (’449 patent) at 8:62–63.

102. I also understand that Patent Owner proposed that “connection properties” means “properties of a connection, including ordering (data transport in order), flow control (a remote buffer is reserved for a connection, and a data producer will be allowed to send data only when it is guaranteed that space is available for the produced data), throughput (a lower bound on throughput is guaranteed), latency (upper bound for latency is guaranteed), the lossiness (dropping of data), transmission termination, transaction completion, data correctness, priority, or data delivery.” EX-1019 (Joint CC and Prehearing Statement) at 6. In my opinion, Goossens2002’s connection properties would thus also meet this proposed construction.

103. **Connection Properties of Channels Are Independently Adjustable.**

In my opinion, a POSITA would have understood that each of the different channels described by Goossens2002 can have different connection properties, *i.e.*, that they

are “adjustable independently.” A POSITA would have understood this because, rather than describing that each *connection* has certain properties, Goossens2002 describes that each *channel* has certain properties. For example, as previously discussed, Goossens2002 explicitly states that “[a] channel can have a combined throughput and latency guarantee (GT), or be best effort (BE), with an unspecified finite latency upper bound.” Ex-1005 (Goossens2002) at 3. Therefore, a POSITA would have understood that in Goossens2002’s NOS, one such channel (*e.g.*, the output channel) may be “best effort (BE), with an unspecified finite latency upper bound,” whereas another (*e.g.*, the input channel) may “have a combined throughput and latency guarantee (GT).”

104. This understanding is further supported by Goossens2003, which describes the same *ÆTHEREAL* NoC as Goossens2002, and says explicitly that the channels can be independently configured to have different properties, *i.e.*, that they are independently adjustable. *See* Ex-1007 (Goossens2003) at 74 (“the request and response communications can be independently configured (*e.g.* for flow control, and throughput guarantees), allowing more fine-grained resource management”).

105. This is also consistent with what a POSITA would have understood about configuring the properties of channels individually to match the corresponding application-level requirements. For example, as I previously discussed in the technology background section, it was well-known in the art that the component

channels of a connection may have different requirements that necessitate different properties, and that appropriate QoS guarantees should be provided for each of these channels. *See, e.g.,* Ex-1012 (Kandlur) at 1045 (“In a bidirectional connection between a pair of user entities, the message generation characteristics may differ substantially for the two directions. Hence, it is preferable to restrict the real-time channel to unidirectional communication, and a bidirectional connection can be composed from a pair of channels.”); Ex-1011 (Campbell) at 10 (“For example, in a videophone connection consisting of separate video and audio channels, a slightly lower commitment probability for the video channel would be appropriate.”).

106. To the extent that Patent Owner or its expert argues that Goossens2002 does not explicitly disclose that its channels are adjustable independently, in my opinion, this would have been obvious in view of Drake. Just like Goossens2002, 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.

107. Drake also describes “AllocationRequest” messages, which are used “[t]o request an allocation of a QoS path for a particular data stream.” Ex-1006 (Drake) at 8:31-32. These messages include, among other fields: (1) a “QoSDesired” field, which indicates “[t]he quality of service desired for this data

stream” and (2) a “QoSAcceptable” field, which indicates “[t]he quality of service acceptable for this data stream.” Ex-1006 (Drake) at 8:43–46. Drake describes that the desired and acceptable QoS are specified using the following “quality-of-service parameters that describe the service expected from the network”:

Throughput—Forward Direction. The average throughput in bits per second over an interval of length T from source station 10 to target station 29.

Throughput—Backward Direction. The average throughput in bits per second over an interval of length T from the target station 29 to the source station 10.

End-to-End Delay—Forward Direction. The maximum elapsed time between the storage of a packet in the LAN entry device at source station 10 and the reception of the packet from the LAN exit device at target station 29.

End-to-End Delay—Backward Direction. The maximum elapsed time between the storage of a packet in the LAN entry device at the target station 29 and the reception of the packet from the LAN exit device at source station 10.

Ex-1006 (Drake) at 5:34–47. These 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. Drake also describes that these parameters can be adjusted after their initial allocation. Ex-1006 (Drake) at 12:35–13:62, Figs. 7–9.

108. Because Drake’s quality-of-services parameters are separately specified for each directional channel, a POSITA would have understood that these channels can be configured to have different properties, such that they are

“adjustable independently.” 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. For example, one amount of throughput (*i.e.*, bandwidth) can be reserved for the forward direction channel and a different amount reserved for the backward direction channel. This is precisely what is disclosed in the ’449 patent, which describes that “the outgoing and return parts of connections can be configured independently” such that “different bandwidths can be reserved for requests and responses.” Ex-1001 (’449 patent) at 11:4–55.

109. These separately configurable channels are consistent with Drake’s description of “traffic descriptors” that include separate fields for the forward direction and backward direction. Ex-1006 (Drake) at 5:51–6:28. These traffic descriptors indicate that the types of communications being transmitted over the two directional channels can differ substantially, necessitating different QoS guarantees. For example, Drake’s “access class” field indicates various types of services that can be provided over one or both of the directional channels, including “access priority, synchronous, asynchronous, etc.” A POSITA would have understood that a channel carrying synchronous data may have a greater need for latency guarantees than a channel carrying asynchronous data.

110. Moreover, Drake also describes different routes or “connection paths” between a source and a destination that can be specified by a “route selection mechanism.” *See, e.g.*, Ex-1006 (Drake) at 2:61–63, 5:8–14. In my opinion, a POSITA would have recognized therefore that each of these connection paths is a channel, and specifically that a two-way connection could include different routes for the forward and backward channels. Additionally, Drake explains the these channels can have different connection properties, such as bandwidth (*i.e.*, throughput) and delay depending on which network components are included in the selected route. Ex-1006 (Drake) at 9:60–62; *see also id.*, at 8:31–35, 10:41–51. Thus, by selecting a route with one set components for the forward channel and a route with a different set of components for the backward channel, the connection properties of the channels can be adjusted independently.

- g. Claim 1[f]: “wherein said connection 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”**

111. In my opinion, this limitation is disclosed by Goossens2002 and also rendered obvious by Goossens2002 in view of Drake.

112. As I previously discussed with respect to claim element 1[e], Goossens2002 explains that its network provides “half duplex and full duplex connections,” *i.e.*, connections that enable communications in two directions. Ex-


1005 (Goossens2002) at 2. Goossens2002 also describes that its connections support “transactions.” Ex-1005 (Goossens2002) at 2. Therefore, in my opinion, a POSITA would have understood that Goossens2002’s connections support outgoing messages from a first IP to a second IP and return messages from the second IP to the first IP. Additionally, as I have shown above, Drake likewise describes communications in the “forward direction” (*i.e.*, outgoing messages) and in the “backward direction” (*i.e.*, return messages).

h. Claim 1[g]: “wherein the resource manager determines whether the at least one connection with said communication channels is available based on the connection properties being available”

113. In my opinion, Goossens2002 in view of Drake renders this limitation obvious. As I previously discussed with respect to claim element 1[c], Goossens2002 in view of Drake discloses a “resource manager” in the form of QoS allocator that receives requests for connections with certain QoS properties. *See* above Claim Element 1[c]. Specifically, the QoS allocator examines the resources of the network in order to determine if there are available resources based on the QoS parameters that are required for the data stream—*i.e.*, whether the “connection properties” are available for the data stream. *See* Ex-1006 (Drake) at Abstract (“An allocator in the management system determines if the requested facilities are available....”), 2:14–21, 2:65–3:6, 4:49–56 (describing that the QoS allocator “examine[s] the resources” of the network and determines if it can “satisf[y] the

requested level of QoS.”), 10:41–51 (QoS allocator “determine[s] if the QoS parameters required for the data stream is available”), 10:52–11:48. Based on the availability of the resources, the QoS allocator determines whether the request connection can be created/established.

114. Drake also explains the QoS allocator’s process for determining the availability of the requested connection properties with reference to Figures 4 and 5. *See* Ex-1006 (Drake) at 10:9–51, 10:52–11:48, Figs. 4–5. Annotated Figure 5 is shown below:



115. First, a requesting module “sends a request to the allocator 26 to determine if the QoS parameters required for the data stream is available.” Ex-1006 (Drake) at 10:41–46. Then, once the QoS allocator receives the request, it “determine[s] if the QoS is available in the components in the multimedia transmission path.” Ex-1006 (Drake) at 10:46–49. Drake explains that this determination is made using a database, which contains information regarding the components of the network and the resources reserved by previously allocated connections. Ex-1006 (Drake) at 5:8–14 (“In carrying out its management functions, manager 23 utilizes the data in QoS allocation data base 25 which, in general, contains all of the quality of service parameters for each of the components of LAN 17, as well as descriptions of all of the data streams already accepted by allocator 20 and the detailed constituents of the connection paths for each of those data streams.”), 10:60–11:48, at 12:11–15 (“It is to be noted that, as described above, each component in the database 25 includes a BandwidthAllocated parameter that represents the amount of bandwidth already allocated to paths in use.”).

116. The QoS allocation database used by Drake’s QoS allocator to determine the availability of requested connection properties is exactly like the “property table” used by the resource manager of the ’449 patent to determine the availability of resources. As the ’449 patent explains, “the resource manager RM may contain a property table with entries for all the different properties of the

channels and connections” and “all different properties of the network.” Ex-1001 (’449 patent) at 10:34–40.

2. Claim 2: “Integrated circuit according to claim 1, further comprising: at least one communication manager (CM) for managing the communication between different modules”


117. Claim 2 depends from claim 1, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 2 is also rendered obvious by the combination of Goossens2002 and Drake.

118. For example, Drake discloses a “communication manager” in the form of a “QoS requestor manager.” Ex-1006 (Drake) at 4:32–37; *see also* Ex-1006 (Drake) at 4:62–65, 5:5–8, Fig. 1. Drake explains there are QoS requestor managers that are associated with source stations, and that the QoS requestor managers manage communications between modules by “assembl[ing] a request to reserve a QoS connection for a multi-media data stream ... between source station 10 and target station 29,” which request is sent to the QoS allocator (*i.e.*, resource manager). *Id.*

119. Indeed, this is consistent with the disclosure of the communication manager described in the specification of the ’449 patent, which explains that “[t]he communication manager CM request[s] a connection with a set of properties between two modules from the resource manager RM.” Ex-1001 (’449 patent) at 10:19–24.

120. Additionally, in my opinion, to the extent Patent Owner argues that the claimed communication manager must have the ability to request that the resource manager release reserved connections and resources (*see* Ex-1006 (Drake) at 10:51–54), that too is disclosed by Drake. *See, e.g.*, Ex-1006 (Drake) at 8:65–9:1 (describing “QoS de-allocation announcement” message), 11:49–12:34 (describing a process for “de-allocating a previously allocated QoS data stream”). Moreover, in my opinion, it would have been obvious to a POSITA that the QoS requestor managers could either accept or reject connections established by the QoS allocator based on the specified QoS parameters.

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:



Ex-1006 (Drake) at Fig. 1 (annotated).

122. Therefore, in my opinion, for the reasons discussed above, a POSITA would have found it obvious to look to Drake's disclosures regarding a QoS requestor manager to obtain the benefits of Drake's known solution to ensuring the availability of network resources for guaranteed services, and would have had a reasonable expectation of success in doing so. Additionally, in my opinion, a POSITA would have been motivated to do so by implementing the functionality of Drake's QoS requestor manager as a circuit on the same integrated circuit chip as Goossens2002's IP blocks and NOS. Indeed, as I described in the background of

the technology section, the hallmark of a system on silicon—like the one described by Goossens2002—is that the components of the system are included on the same chip, and the nature of Drake’s QoS requestor managers—which intermediate communications between functional blocks over the network—require that they be tightly integrated with Goossens2002 IP blocks and QoS manager. *See* above Section V.A. As such, in my opinion, Goossens2002 in view of Drake renders obvious a “circuit that manages communication between modules,” which I understand is the construction of “communication manager” that patent owner has proposed. Ex-1019, at 12.

3. Claim 3: “The Integrated circuit of claim 1, further comprising at least one of a switch and a router”

123. Claim 3 depends from claim 1, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 3 is also rendered obvious by the combination of Goossens2002 and Drake, as Goossens2002 describes that its system includes both routers and switches.

124. For example, as I previously discussed with respect to claim element 1[b], Goossens2002 explains that its NoS contains “a number of routers.” *See* above Claim Element 1[b]; Ex-1005 (Goossens2002) at 2. Additionally, Goossens2002 further describes that its NOS contains switches to implement a form of circuit switching. Ex-1005 (Goossens2002) at 3 (“We use a combination of circuit and packet switching, in the spirit of ATM [1].” “GT streams use time-division-

multiplexed circuit switching.”). Therefore, to the extent the claim requires “at least one switch and at least one router” (Ex-1019 at 7), in my opinion, this is disclosed by Goossens2002.

4. Claim 4: “The Integrated circuit of claim 1, further comprising a chip, wherein the processing modules and the network are disposed on said chip”

125. Claim 4 depends from claim 1, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 4 is also rendered obvious by the combination of Goossens2002 and Drake for the reasons that I explained previously with respect to claim element 1[d], which includes all of the same requirements. *See above* Claim 1[d]. In particular, claim element 1[d] requires that the processing modules and network “are disposed on a chip.”

5. Claim 5: “The Integrated circuit of claim 1, wherein the connection properties comprise at least one of data transport ordering, flow control, throughput, latency, and lossiness”

126. As I previously discussed, in my opinion, Claim 1 is obvious as discussed above. Additionally, in my opinion, as I previously discussed with respect to claim element 1[e], Goossens2002 describes that each of its channels has connection properties that include throughput and latency. Ex-1005 (Goossens2002) at 2–3 (“A channel can have a combined throughput and latency guarantee (GT)”). Drake likewise describes that its channels have properties that

include “throughput” and “end-to-end delay” (*i.e.*, latency). Ex-1006 (Drake) at 5:30–50.

6. Claim 6

a. Claim 6[pre]: “Integrated circuit according to claim 2”

127. As I previously discussed, in my opinion, Claim 2 is obvious as discussed above. *See* above Claim 2.


b. Claim 6[a]: “wherein said first module (M; I) is adapted to issue a request (REQ) for a connection with at least one of said second modules to said communication manager (CM)”

128. In my opinion, Goossens2002 in view of Drake renders this limitation obvious. For example, Goossens2002 describes that its IPs (*i.e.*, processing modules) “can set up and remove connections.” Ex-1005 (Goossens2002) at 3. Goossens2002 explains that the IPs do so by making a “service request” that includes information regarding the performance requirements of the connection. Ex-1005 (Goossens2002) at 2–3. The IP that makes a service request is a “first module” and the IP that is the destination or target of the connection is a “second module.”

129. Similarly, Drake describes that an application in a source station (*i.e.*, first module) “requests a QoS connection to the destination or target station 29” (*i.e.*, second module). Ex-1006 (Drake) at 4:26–29; *see also* Ex-1006 (Drake) at Abstract, 5:25–50, 2:32–41 (describing “requested QoS values”). In my opinion, a POSITA

would have understood from these disclosures that the requests described by Goossens2002 and Drake include desired connection properties.

130. Additionally, Drake also describes that the application in the source station (*i.e.*, first module) communicates this request to a QoS requestor manager (*i.e.*, communication manager) for subsequent transmission to a QoS allocator, as discussed further below. *See* Ex-1006 (Drake) at 3:21–28, 4:32–37, 4:62–65 (“The QoS requestor manager 12 interacts with application using internal messages as is well known in the art”), 5:5–8, Fig. 1. This is indicated by the arrows shown in Drake’s Figure 1 annotated below:



Ex-1006 (Drake) at Fig 1 (excerpted, annotated).

- c. **Claim 6[b]: “said communication manager (CM) is adapted to forward said request (REQ) for a connection with communication channels each having a specific set of connection properties to said resource manager (RM)”**


131. In my opinion, Goossens2002 in view of Drake renders this limitation obvious. For example, Drake describes that in having received a request for a connection from an application in the source station (*i.e.*, first module), the QoS requestor manager (*i.e.*, communication manager) “assemble[s] a request to reserve a QoS connection.” Ex-1006 (Drake) at 4:32–37. Drake also explains that the QoS requestor manager “utilizes QoS protocol machine 13 to format the QoS request,” which is also referred to by Drake as a “QoS allocation request.” Ex-1006 (Drake) at 4:62–65; *see also id.* at 8:31–49, 9:52–10:8 (describing that the QoS request “contains a specification of the quality of service which is desired” and the minimum acceptable quality of service), Fig. 3. Thus, the QoS request is then sent to and received by the QoS allocator (*i.e.*, resource manager). Ex-1006 (Drake) at 4:50–54 (“The overall function of allocator 20 is to receive the QoS request from station 10....”); *see also* Ex-1006 (Drake) at 5:5–8 (“Allocator station 20 ... receives, decodes and acts upon the QoS request from source station 10.”).

132. In addition, as I previously discussed above with respect to claim element 1[e], the connections requested by Drake’s QoS requestor manager

comprise communication channels each having a specific set of connection properties. *See* above Claim Element 1[e].

d. Claim 6[c]: “said resource manager (RM) responding with the availability of the requested connection to said communication manager (CM)”

133. In my opinion, Goossens2002 in view of Drake renders this limitation obvious. As I previously discussed with respect to claim element 1[g], Drake describes that, when the QoS allocator receives requests for connections with certain QoS properties, it examines the resources of the network to “determine if the QoS parameters required for the data stream is available,” *i.e.*, whether a connection with the requested connection properties is available. *See* above Claim Element 1[g]. Additionally, Drake further describes that, once the QoS allocator has made its determination, it sends a “QoS allocation reply,” which indicates, among other things, whether the requested connection and connection properties are available. Ex. 1006 (Drake) at 4:55–57 (describing a reply sent if “the QoS is not available,”), 8:20–29 (describing “QoS allocation reply” message), 8:50–64 (describing contents of “AllocationReply” message including “QoS granted”); Fig. 5. This is shown below in annotated Figure 5:



Ex. 1006 (Drake) at Fig. 5 (annotated).

- e. **Claim 6[d]: “wherein a connection between the first and second module is established based on the available properties of said communication channels of said connection”**

134. In my opinion, Goossens2002 in view of Drake renders this limitation obvious. For example, Drake describes that if the QoS allocator determines that the

requested QoS properties are available, the requested connection between the source and target stations is established along a designated path through the network and required resources are reserved. *See, e.g.*, Ex-1006 (Drake) at Abstract (“An allocator in the management system determines if the requested facilities are available and, if so, reserves these facilities for the requesting station.”); 10:9–11 (describing “the process of allocating a QoS transmission path between a source of multimedia data and a user of such data.”), 11:17–19, 11:26–28 (“If the allocation has been accepted, ...the multimedia file is transmitted to the target station 29”), Figs. 4–5. This newly established connection is identified and described in the “QoS allocation reply” message by a “StreamID,” “Route Designator Fields,” and an identification of the “QoS parameters of the path granted” to the connection. *See* Ex-1006 (Drake) at 8:50–64. In addition, Drake describes that the database is updated to reflect the resources that have been reserved for the newly established connection, and are therefore unavailable for future connection allocations. *See* Ex-1006 (Drake) at 11:11–17 (describing that reserved bandwidth is “subtracted from the previously available bandwidth on the components in the multimedia data path to determine the new bandwidth available for the next allocation”), Fig. 5.

7. Claim 9: “Integrated circuit according to claim 2, further comprising: at least one network interface means (NI), associated to each of said modules, for managing the communication between said modules and said network (N)”

135. Claim 9 depends from claim 2, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 9 is also rendered obvious by the combination of Goossens2002 and Drake.

136. Goossens2002 describes “network interfaces” that “implement multipath and multi-channel ordering, and channel flow control.” Ex-1005 (Goossens2002) at 3; *see also* Ex-1005 (Goossens2002) at 3 (describing that channels implement “point-to-point communication between two network interfaces”). Goossens2002 further describes that the computation performed by IPs is decoupled from communication over the NOS. Ex-1005 (Goossens2002) at 1. Therefore, in my opinion, a POSITA would have known that Goossens2002’s system on silicon includes network interfaces between the IP modules and the NOS, and that the network interfaces manage the communication between the IPs and the NoS.

8. Claim 10

137. In my opinion, the combination of Goossens2002 and Drake teaches all the limitations of claim 10, and therefore would have rendered that claim obvious as I discuss below.

- a. **Claim 10[pre.1]: “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”**

138. In my opinion, to the extent it is limiting, this part of the preamble is disclosed by Goossens2002. As I previously discussed with respect to claim elements 1[pre]–1[b] and 1[d], Goossens2002 describes a system on silicon (*i.e.*, integrated circuit) in which multiple intellectual property blocks (*i.e.*, modules) communicate with one another over connections via a “network on silicon.” *See* above Claim Elements 1[pre]-1[b] and 1[d]; *see also* Ex-1005 (Goossens2002) at 1 (describing “a network on silicon (NOS) ... to implement communication between IP cores”), 3 (describing that “IPs can set up and remove connections”). Additionally, Goossens2002 further describes that these IPs “interact in many different ways” including “message passing.” Ex-1005 (Goossens2002) at 2. These messages are communicated over the connections in the form of packets. *See* Ex-1005 (Goossens2002) at 1 (describing a “packet-switched router network”), 2 (describing “a NOS based on packet-switching”), 3 (similar).

- b. **Claim 10[pre.2]: “wherein said connections comprises a set of communication channels each having a set of connection properties, any communication channel being independently configurable”**

139. In my opinion, to the extent it is limiting, this part of the preamble is substantively the same as element 1[f] of claim 1, and is disclosed by Goossens2002

and/or rendered obvious by Goossens2002 in view of Drake for the same reasons that I explained with respect to that claim element. *See* above Claim Element 1[e]. Additionally, I understand that Patent Owner proposed a construction of “independently configurable” to mean “able to be configured with connection properties that do not depend on those of another configuration.” EX-1030 at 13. In my opinion, Goossens2002 alone or in view of Drake also renders this limitation obvious under this construction.

- c. **Claim 10[pre.3]: “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”**

140. In my opinion, to the extent it is limiting, this part of the preamble is substantively the same as element 1[f] of claim 1, and is disclosed by Goossens2002 for the same reasons that I explained with respect to that claim element. *See* above Claim Element 1[f].

- d. **Claim 10[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”**

141. In my opinion, this limitation is substantively the same as element 6[a] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See* above Claim Element 6[a].

e. Claim 10[b]: “the communication manager forwarding the request to a resource manager”

142. In my opinion, this limitation is substantively the same as element 6[b] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See* above Claim Element 6[b].

f. Claim 10[c]: “the resource manager determining whether a target connection with the desired connection properties is available”

143. In my opinion, this limitation is substantively the same as element 1[g] of claim 1, and is obvious for the same reasons that I explained with respect to that claim element. *See* above Claim Element 1[g].

g. Claim 10[d]: “the resource manager responding with the availability of the target connection to the communication manager”

144. In my opinion, this limitation is substantively the same as element 6[c] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See* above Claim Element 6[c].

h. Claim 10[e]: “the target connection between the first and second module being established based on the available properties of said communication channels of said connection”

145. In my opinion, this limitation is substantively the same as element 6[d] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See* above Claim Element 6[d].

9. Claim 11: “The method of claim 10, wherein the network manages traffic utilizing at least one of a switch and a router”

146. Claim 11 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 11 is also rendered obvious by the combination of Goossens2002 and Drake for the reasons that I explained previously with respect to claim 3, which includes an identical limitation. *See above Claim 3.*

10. Claim 12: “The method of claim 10, wherein the processing modules and the network are disposed on a chip”

147. Claim 12 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 12 is also rendered obvious by the combination of Goossens2002 and Drake for the reasons that I explained previously with respect to claim 4, which includes an identical limitation. *See above Claim 4.*

11. Claim 13: The method of claim 10, wherein the connection properties comprise at least one of data transport ordering, flow control, throughput, latency, and lossiness”

148. Claim 13 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 13 is also rendered obvious by the combination of Goossens2002 and Drake for the reasons that I explained previously with respect to claim 5, which includes an identical limitation. *See above Claim 5.* I also understand that Patent Owner proposed a


construction that “flow control” means “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-1019, at 6. In my opinion, Goossens2002 in view of Drake renders this limitation obvious under this construction. *See* EX-1005 (Goossens2002) at 2 (“Flow control avoids data loss by aligning IP production and consumption rates.”).

12. Claim 14: “The method of claim 10, further comprising the resource manager determining whether the target connection with the desired connection properties is available based on reading of a centralized or distributed property table comprising properties associated with the network”

149. Claim 14 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 14 is also rendered obvious by the combination of Goossens2002 and Drake

150. As I previously discussed with respect to claim element 1[g], Drake describes that its QoS allocator (*i.e.*, resource manager) determines whether a requested connection is available by referencing a “QoS allocation data base,” which a POSITA would recognize as a centralized property table. Drake’s QoS allocation data base contains information regarding the properties associated with the components of the network and the resources reserved by previously allocated connections. *See* above Claim Element 1[g]; *see also* Ex-1006 (Drake) at 3:14–20 (“network topology information and current data stream allocations are maintained

in a hierarchical database”), 5:14–8:19, (describing “the detailed contents of the database 25”), Fig. 2. A POSITA would have understood that databases are typically constituted in a table-format. For example, Drake’s QoS allocator queries the database to determine whether requested connection properties, such as “bandwidth” (i.e., throughput) and/or “maximum delay” (i.e., latency), are available, and if so, the requested connection is established and the required resources are reserved. Ex-1006 (Drake) at 10:63–11:7; *see also* Abstract. The QoS allocation data base is shown below in Drake’s annotated Figure 1.



Ex-1006 (Drake) at Fig. 1 (annotated).

13. Claim 15: “The method of claim 10, further comprising the resource manager reserving the target connection after determining whether the target connection with the desired connection properties is available”

151. Claim 15 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 15 is also rendered obvious by the combination of Goossens2002 and Drake.

152. As I previously discussed with respect to claim element 1[g], Drake describes that, when the QoS allocator (*i.e.*, resource manager) determines that a requested connection is available, the QoS allocator reserves the resources required for the requested connection. *See* above Claim 1[g]; *see also, e.g.*, Ex-1006 (Drake) at Abstract (“An allocator in the management system determines if the requested facilities are available and, if so, reserves these facilities for the requesting station.”); 4:31–36 (describing a request to “a request to reserve a QoS connection”), 7:45–8:5 (“QoSStream 42–Represents a QoS reserved end-to-end data stream”).

14. Claim 16: “The method of claim 10, wherein the connection properties comprise at least one of throughput, latency and jitter”

153. Claim 16 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 16 is also rendered obvious by the combination of Goossens2002 and Drake for the reasons described with respect to claim 5. *See* above Claim 5.

154. As I explained with respect to that claim, Goossens2002 describes that each of its channels has connection properties that include throughput and latency. Ex-1005 (Goossens2002) at 2–3 (“A channel can have a combined throughput and latency guarantee (GT)”). Drake likewise describes that its channels have properties that include “throughput” and “end-to-end delay” (*i.e.*, latency). Ex-1006 (Drake) at 5:30–50.

B. Ground II: Claims 1-6 and 9-16 Are Rendered Obvious by Goossens2003 in view of Drake

155. In my opinion, claims 1–6 and 9–6 are rendered obvious by the combination of Goossens2003 and Drake.

156. As a threshold matter, a POSITA would have been motivated to combine the disclosures of these two references because both relate to the use of “connections” to provide guaranteed services, and Drake provides additional implementation details that would aid a POSITA in creating an SoC with improved performance and efficiency.

157. As I summarized above, Goossens2003 describes a network on chip (NOC) architecture for SoCs. Ex-1007 (Goossens2003) at Abstract, 63, 74. Goossens2003 describes that a feature of this NoC is the provision of guaranteed services to support real time and critical functions. Ex-1007 (Goossens2003) at 71 (“In this section we introduce the *ÆTHEREAL* NOC. It offers guaranteed services to obtain the advantages in composability, robustness, and so on of QOS-based

design. These services are used for real-time and critical functions.”). The paper describes that the available services pertain to data correctness on delivery, completion of data delivery, and bounds on performance that include completion time, latency, jitter, and energy consumption. Ex-1007 (Goossens2003) at 73.

158. Goossens2003 describes that these guaranteed services are provided via “connections.” Ex-1007 (Goossens2003) at 71. As I explained in the technology background section above, by the early 2000’s, connections were a well-known mechanism for providing guaranteed services. Goossens2003 further describes a negotiation phase in which connections are requested with explicitly defined service levels, which are either accepted or rejected. Ex-1007 (Goossens2003) at 71.

159. However, Goossens2003 does not include certain details regarding the implementation of its guaranteed services and connections. For example, the paper explains that requests for connections will be granted if there are sufficient resources available, and that they will otherwise be rejected, but does not describe a component that makes this determination or the process by which it does so. Ex-1007 (Goossens2003) at 67. Goossens2003 also describes that the establishment of connections to provide guaranteed services may require the reservation of network resources, which must be carefully managed given the limited resources in SoCs. Ex-1007 (Goossens2003) at 65, 67. Goossens2003, however, does not describe a

component or mechanism for managing resource usage within the NOC or for determining the availability of the resources required for a requested connection.

160. As I explained in Ground I, Drake provides details regarding the implementation of these features and a POSITA would have found Drake's known solution for "ensur[ing] selected levels of quality of services" helpful in implementing an SoC like the one described by Goossens2003. *See* Ex-1006 (Drake) at Abstract. A POSITA would have been particularly interested in Drake's disclosure of a "QoS requestor manager" to formulate requests for guaranteed services with particular properties and a "QoS allocator" to determine whether the network resources required for the requested services are available. *See* Ex-1006 (Drake) at 4:17–56; *see also* Ex-1006 (Drake) at 2:14–40.

161. Like Goossens2003, Drake describes providing guaranteed services for certain kinds of network communications, including guarantees related to throughput and latency. Drake describes that these services are particularly advantageous for multimedia applications, which require large amounts of bandwidth to transmit video data and bounded delay in the delivery of the data. Ex-1006 (Drake) at 1:22–46. Also, like Goossens2003, Drake describes providing these guaranteed services using "connections." *See, e.g.*, Ex-1006 (Drake) at 4:26–29.

162. Drake recognizes that a critical element in providing such guaranteed services is ensuring the availability of the required network resources. To do so,

Drake describes a quality of service management system that includes a “QoS requestor manager” to formulate requests for guaranteed services with particular properties and a “QoS allocator” to determine whether the network resources required for the requested services are available. Ex-1006 (Drake) at Abstract, 2:14–40, 4:17–56, Figs. 1, 4–5. When the required resources are available, they are reserved for the pending request in order to ensure their availability, and when they are not available the request is denied. Ex-1006 (Drake) at Abstract. Drake explains that “[t]he intent of such a technique is to deny network access to those streams whose level of quality of service requirements are not currently available.” Ex-1006 (Drake) at 2:18–21.

163. In my opinion, a POSITA would have understood that using Drake’s quality of service management process would improve the efficiency and performance of Goossens2003’s network on chip by ensuring that the network resources required to provide guaranteed services are actually available, and that the network will be able to fulfill its commitments to the IPs that have been granted such services. Indeed, Goossens2003 describes the same goal, explaining that “The NOC that offers communication services will honour the request if has sufficient resources available, and will reject it otherwise.” Ex-1007 (Goossens2003) at 67.

164. In my opinion, a POSITA would also have reasonably expected a combination of the disclosures of Goossens2003 and Drake to be successful. For

instance, a POSITA would have recognized that these two references are both in the same field of study—guaranteed services for data communications in packet-based networks—and both are directed to ensuring that sufficient resources are available for the communication of multimedia content, such as audio and video. *See, e.g.*, Ex-1006 (Drake) at 1:22–24 (describing “audio and full motion video”); Ex-1007 (Goossens2003) at 62 (describing quality of service for “real-time applications ... involving audio and video”). A POSITA would also have recognized that such a combination would have involved nothing more than combining prior art elements according to known methods to yield predictable results. Drake provides details about a system for providing guaranteed services that is consistent with and would have worked with the system of Goossens2003. Therefore, in my opinion, a POSITA would have reasonably expected that use of Drake’s known solution to ensuring the availability of network resources for guaranteed services would improve Goossens2003’s NOC in the same way.

165. In addition, while Drake describes its QoS management system in the context of local area networks, in my opinion, a POSITA would recognized that the advantages and benefits of Drake’s system are applicable to a network on silicon like the one described by Goossens2003. For instance, Goossens2003 itself notes that certain resources are more limited in the on-chip context and need to be managed carefully. *See* Ex-1007 (Goossens2003) at 65 (“SOCs ... have limited resources,

and they must be shared and managed.”). Drake’s QoS management process serves this very purpose and its more efficient utilization of network resources is well-suited for such a resource-constrained environment.

166. Furthermore, the parallels between mature off-chip networks and newer on-chip networks would have been well-known to POSITA. As I explained above, NoCs have borrowed many design features from scalable interconnection networks that were designed for large-scale parallel machines in the 1990s. For example, as described in one text devoted to networks-on-chip “[d]esigning a network for on-chip communication is similar to the problem of designing a local area computer network with some computing and communication requirements.” Ex-1014 (Jantsch) at 102. Another chapter in the same text explained that addressing issues related to network resource allocation—as Goossens²⁰⁰² and Drake do—is substantially similar in on-chip and off-chip networks: “The resource allocation, optimization of network usage and verification of performance and correctness are the main problems that are basically similar to what distributed and parallel system designers have to face.” Ex-1014 (Jantsch) at 33. Indeed, the ’449 patent itself acknowledges that many of the issues faced by networks on chip had been studied in the context of off-chip networks and that results of those studies are applicable to on-chip networks as well. Ex-1001 (’449 patent) at 1:50–66.

1. Claim 1

167. In my opinion, Goossens2003 in view of Drake teaches all the limitations of claim 1, and thus would have rendered that claim obvious.

a. Claim 1[pre]: “Integrated circuit comprising”

168. In my opinion, to the extent the preamble is limiting, Goossens2003 discloses an integrated circuit. Specifically, Goossens2003 is directed to communications in “systems on chip (SOC).” *See, e.g.*, Ex-1007 (Goossens2003) at 61 (describing “quality of service (QoS) ... in the context of ... systems on a chip (SOC)”); *see also id.*, Abstract, 63–65, 67, 68, 71, 78. As I described above in the background of the technology (*see* above Section V.A) a POSITA would have understood that an SOC is implemented on a single integrated circuit chip, and thus all of the components of the system described by Goossens2003 are on the same chip. *See, e.g.*, Ex-1014 (Jantsch) at vii (“During the 1990s more and more processor cores and large reusable components have been integrated on a single silicon die, which has become known under the label System on Chip (SoC).”); Ex-1016 (Man) at 12 (“A SOC thus encapsulates system knowledge in a single medium—silicon—which blurs the separation between component design and system design.”); Ex-1008 (Wolf) at 1 (“An SoC is an integrated circuit that implements most or all of the functions of a complete electronic system.”); Ex-1009 (Badawy) at 1 (“With the increase in the number of logic gates that can be implemented on a single chip,

various functionalities ... can be integrated into a single silicon chip, realizing an entire system on chip (SoC)"). In my opinion, a POSITA would have understood that a system *on chip* is implemented on a single integrated circuit chip.

169. Indeed, in discussing such SOC's, Goossens2003 describes "the number of transistors that can be integrated in an IC," or integrated circuit, and "the intrinsic computational power of an IC." Ex-1007 (Goossens2003) at 63. Additionally, as I discuss below with respect to claim elements 1[a]–[b], Goossens2003's SOC includes IP blocks interconnected by a network on chip. *See* Below Claim Elements 1[a]–[b]. Therefore, in my opinion, Goossens2003 discloses an integrated circuit, including under the construction that I understand Patent Owner has proposed in parallel district court proceedings: "interconnected circuitry on a chip." *See* Ex-1019 at 9.

b. Claim 1[a]: "a plurality of processing modules (M; I; S; T)"

170. In my opinion, Goossens2003 discloses this limitation. Specifically, Goossens2003 describes that systems on chip include multiple "intellectual property blocks (IP)," which it also refers to as "masters" and "slaves." Ex-1007 (Goossens2003) at 63; *see also, e.g.,* Ex-1007 (Goossens2003) at 61, 67–69 (describing "combining IPs to obtain a SOC"), 71–74 (describing "communication[s] between one master and one or more slaves"), 77–78. In my opinion, a POSITA would have understood that these IPs constitute processing

modules. For example, Figure 4.5 of Goossens2003, which I have annotated below, depicts a number of such IPs, including one labeled “ μ P”—microprocessor—and two labeled “M,” for “memory module.”





Figure 4.5. A narrowcast connection.

Ex-1007 (Goossens2003) at Fig. 4.5 (annotated); *see also* Ex-1007 (Goossens2003) at 72. In my opinion, this is also consistent with the specification of the '449 patent, which describes that the “modules” of its disclosed embodiments “can be so-called intellectual property blocks IPs (computation elements, memories or a subsystem which may internally contain interconnect modules).” Ex-1001 ('449 patent) at 9:42–46. Indeed, the '449 patent contains a nearly identical figure.



Ex-1001 ('449 patent) at Fig. 1

c. Claim 1[b]: “network (N; RN)”

171. In my opinion, Goossens2003 discloses this limitation. Specifically, Goossens2003 explains that IPs communicate via “networks on chip (NOC).” *See, e.g.,* Ex-1007 (Goossens2003) at Abstract, 64, 66 (describing “[u]sing a NOC for the platform interconnect” to “integrate[] heterogeneous IPs in a standard fashion”), 68–69 (describing “[a]n IP connected to a NOC”), 71–77. Goossens2003 also explains that this NOC is comprised of “routers” that “transport data.” Ex-1007 (Goossens2003) at 74. Goossens2003 illustrates a portion of its NOC—including three routers—in Figure 4.6, below below:




Figure 4.6. Three routers, all in slot $s = 3$, switch packets p_i of circuit c_i to outputs o_i , for $T(s, o_i) = c_i$. For each table, only the relevant columns (outputs) are shown. Input/output link pairs are labelled io_i .

Ex-1007 (Goossens2003) at Fig. 4.6.

172. I understand that Patent Owner has advocated in parallel district court proceedings for a construction of “network” as “a multi-hop interconnection between modules separated by one or more network nodes.” Ex-1019 at 9. To the extent Patent Owner urges such a construction here, in my opinion, a POSITA would

have understood that Goossens2003 discloses such an interconnect means because its routers are “network nodes” and the need for a packet to traverse one or more of such routers from its source to its destination means it includes a “multi-hop interconnection between modules separated by one or more network nodes.” *See* Ex-1019 at 2 (“Each interconnection between routers may be called a ‘hop,’ such that the NoC is a ‘multi-hop’ interconnection between modules separated by routers, or network nodes.”). As described above with reference to Figure 4.6, Goossens2003 discloses multi-router connections, such that a path can have multiple hops (through router 1, 2, and 3). A POSITA would understand these routers to be “network nodes.” *See, e.g.*, Ex-1023 (Kanada) at 161 (“However, we now need programs for configuration because the node functions to be configured are much more complex in today’s network nodes, e.g., QoS-ready routers.”). Indeed, Goossens2003 distinguishes its NOC from “single-hop on-chip communication such as busses and switches.” Ex-1007 at 77.

d. Claim 1[c]: “at least one resource manager (RM) for managing resources of the network (N)”

173. In my opinion, Goossens2003 in view of Drake renders this limitation obvious. As I discussed above, Goossens2003 explains that the implementation of connections in a NOC requires “resource management,” but 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.

See also Ex-1007 (Goossens2003) at 65, 68–69, 74. As I previously discussed above in Ground I, however, Drake discloses a “resource manager” in the form of “a quality of service (QoS) allocator.” For the reasons I previously discussed above, in my opinion, a POSITA would have looked to Drake’s disclosures regarding a QoS allocator for details regarding the implementation of resource management in Goossens2003’s NOC, and would have had a reasonable expectation of success in doing so. Additionally, as I discussed above, in my opinion, a POSITA would have looked to Drake’s disclosures regarding a QoS allocator for details regarding the implementation of resource management in Goossens2003’s NOC, and would have been motivated to implement the functionality of Drake’s QoS allocator in a circuit on the same integrated circuit chip as Goossens2003’s IP blocks and NOC. As I previously explained in the background of the technology, the hallmark of a system on chip—like the one described by Goossens2003—is that the components of the system are included on the same chip, and the nature of Drake’s QoS allocator—which manages the resources of the network and intermediates communications between functional blocks—requires that it be tightly integrated with those components. *See* above Section V.A. As such, in my opinion, Goossens2003 in view of Drake renders obvious a “resource manager,” including under Patent Owner’s proposed construction of “circuit that manages network resources.”

- e. **Claim 1[d]: “wherein the modules, network and resource manager are disposed on a chip and arranged for providing at least one connection between a first and at least one second module”**

174. In my opinion, Goossens2003 in view of Drake renders this limitation obvious. As I previously discussed above with respect to claim elements 1[pre]–1[b], Goossens2003 describes a system on chip that includes multiple IPs and a NOC, and a POSITA would have understood that all of these components are on the same chip. *See* above Claim Elements 1[pre]-1[b].

175. Additionally, as I previously discussed (*see* above Claim Element 1[c]) a POSITA would have been motivated to include a circuit implementing the functionality of Drake’s QoS allocator on the same chip.

176. Additionally, in my opinion, Goossens2003 discloses that its network on chip is arranged to provide “connections” between the various IPs. For example, Goossens2003 describes that its NOC “offers differentiated services with *connections*” and explains these connections are between a “master” (*i.e.*, first module) and “one or more slaves” (*i.e.*, second modules). Ex-1007 (Goossens2003) at 71; *see also* Ex-1007 (Goossens2003) at 67, 72 (describing “simple,” “narrowcast,” and “multicast” connections). These connections are illustrated below in Goossens2003 Figures 4.4 and 4.5.




Figure 4.4. Connection types.




Figure 4.5. A narrowcast connection.

Ex-1007 (Goossens2003) at Figs. 4.4–4.5; *see also* Ex-1007 (Goossens2003) at 71 (“Different connections are created and closed independently, possibly at different points in time.”) and 67 (“The NOC will release the resources of the connection for use by other connections, when it is closed by the microprocessor.”).

177. As I previously discussed above in Ground I, in my opinion, Drake likewise describes “connections” between an application in a source station (*i.e.*, first module) and a target station (*i.e.*, second module).

- f. **Claim 1[e]: “wherein the at least one connection comprises a set of communication channels each having a set of connection properties, the connection properties of the different communication channels of said connection being adjustable independently”**

178. In my opinion, Goossens2003 discloses this limitation. Specifically, Goossens2003 describes that its connections comprise a “request” channel and a “response” channel, and explains that these channels have “connection properties,” including flow control and throughput guarantees, that “can be independently

configured.” Ex-1007 (Goossens2003) at 74 (“the request and response communications can be independently configured (e.g. for flow control, and throughput guarantees), allowing more fine-grained resource management”); *see also* Ex-1007 (Goossens2003) at 72–74 (describing “connection properties”). Indeed, this is the same functionality that is disclosed in the ’449 patent, which describes that “the connection properties of the request and response parts of a connection can be configured independently for all of throughput, latency and jitter.” Ex-1001 (’449 patent) at 11:4–55. In my opinion, Goossens2003 therefore discloses the claimed connection properties, including under Patent Owner’s proposed construction. Ex-1019, at 6.

179. Additionally, as I discussed above in Ground I, 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. *See* above. In my opinion, these are the same channels and properties disclosed by Goossens2003.

g. Claim 1[f]: “wherein said connection 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”

180. In my opinion, Goossens2003 discloses this limitation. Specifically, Goossens2003 describes that “[o]nce a connection has been created, the master

initiates *transactions* by means of *requests* which zero or more slaves execute, perhaps leading to a *response*.” Ex-1007 (Goossens2003) at 72 (“The master initiates transaction which the slave executes, possibly resulting in a response....”); *see also, e.g.*, Ex-1007 (Goossens2003) at 73 (“Within a single transaction ... the master sends a request, a slave receives the request and executes it, the slave sends the response, the master receives the response.”).

181. As I previously discussed above in Ground I, Drake similarly describes communications in the “forward direction” (*i.e.*, outgoing messages) and in the “backward direction” (*i.e.*, return messages). *See* above.

h. Claim 1[g]: “wherein the resource manager determines whether the at least one connection with said communication channels is available based on the connection properties being available”

182. In my opinion, Goossens2003 in view of Drake renders this limitation obvious. For example, Goossens2003 describes that when requests for connections are issued by IPs, its NOC “will honour the request if has sufficient resources available, and will reject it otherwise.” Ex-1007 (Goossens2003) at 67. But Goossens2003 does not describe the component of its NOC that determines whether sufficient resources are available, nor the process by which it does so.

183. As I previously discussed above in Ground I, however, Drake explains that its QoS allocator (*i.e.*, resource manager) receives requests for connections with certain QoS properties and “determine[s] if the QoS parameters required for the data

stream is available” such that the request can be satisfied and the connection established. *See above*. Therefore, in my opinion, for the reasons I discussed above, a POSITA would have looked to Drake’s disclosures regarding a QoS allocator for details regarding the implementation of resource management in Goossens2003’s NOC, and would have had a reasonable expectation of success in doing so. *See above*.

2. Claim 2: “Integrated circuit according to claim 1, further comprising: at least one communication manager (CM) for managing the communication between different modules”

184. Claim 2 depends from claim 1, which is obvious for the reasons that I explained above. Additionally, in my opinion, Goossens2003 in view of Drake renders this limitation obvious.

185. For example, Goossens2003 describes that its system includes “network interfaces,” and explains that that “end-to-end” connection properties—those properties which relate to the interactions between the master and slave modules, including “reordering, transaction completion, and flow control”—“are implemented by the network interfaces.” Ex-1007 (Goossens2003) at 76–77. Additionally, Goossens2003 further describes that “IPs negotiate with network interfaces to obtain connections with certain properties.” Ex-1007 (Goossens2003) at 77. In my opinion, this is just like the interaction between the modules and

communication manager described in the '449 patent. Ex-1001 ('449 patent) at 10:3–54.

186. Therefore, in my opinion, a POSITA would have understood that Goossens2003's network interfaces manage the communication between different modules. Additionally, in my opinion, Goossens2003's interface discloses the communication manager under Patent Owner's proposed construction. Ex-1019, at 12. Indeed, as previously discussed Goossens2003 discloses that its system is an integrated circuit (*see* Claim Element 1[pre]) and therefore, a POSITA would understand that the Goossens2003's network interfaces are circuits implemented on the system integrated circuit.

187. Additionally, as I previously discussed in Ground I, Drake discloses a “communication manager” in the form of a “QoS requestor manager.” Thus, for the reasons I discussed above, in my opinion, a POSITA would have looked to Drake's disclosures regarding a QoS requestor manager to obtain the benefits of Drake's known solution to ensuring the availability of network resources for guaranteed services and would have had a reasonable expectation of success in doing so.

3. Claim 3: “The Integrated circuit of claim 1, further comprising at least one of a switch and a router”

188. Claim 3 depends from claim 1, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 3 is also rendered obvious by the combination of Goossens2003 and Drake.

189. For example, Goossens2003 describes that its NOC utilizes both switches and routers. As I previously discussed above with respect to claim element 1[b], Goossens2003 describes that its NOC is comprised of “routers” that “transport data.” See above Claim Element 1[b]. These routers are shown below in Figure 4.7 of Goossens2003, which I have annotated:




Figure 4.7. Two views of the combined GT-BE router.

Ex-1007 (Goossens2003) at Fig. 4.7 (annotated). Goossens2003 explicitly states: “[t]he logically separate guaranteed (GT) and best-effort (BE) routers are combined (Figure 4.7(a)) to share the router resources (*e.g. switch and data path*, Figure 4.7(b)), and to obtain the advantages of both.” Ex-1007 (Goossens2003) at 76. Additionally, Goossens2003 further describes that its NOC utilizes a switch to control whether “guaranteed (GT)” or “best-effort (BE)” packets have access to the data path. Ex-1007 (Goossens2003) at 75–76. Thus, to the extent the claim requires

“at least one switch and at least one router” (Ex-1019 at 7), in my opinion, this is disclosed by Goossens2003.

4. Claim 4: “The Integrated circuit of claim 1, further comprising a chip, wherein the processing modules and the network are disposed on said chip”

190. Claim 4 depends from claim 1, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 4 is also rendered obvious by the combination of Goossens2003 and Drake for the reasons that I explained previously with respect to claim element 1[d], which includes all of the same requirements. *See* above Claim 1[d]. In particular, claim element 1[d] requires that the processing modules and network “are disposed on a chip.”

5. Claim 5: “The Integrated circuit of claim 1, wherein the connection properties comprise at least one of data transport ordering, flow control, throughput, latency, and lossiness”

191. Claim 5 depends from claim 1, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 5 is also rendered obvious by the combination of Goossens2003 and Drake. As I previously discussed above with respect to claim element 1[e], Goossens2003 describes that each of its channels has properties that include at least “flow control” and “throughput guarantees.” Ex-1007 (Goossens2003) at 74; *see also* Ex-1007 (Goossens2003) at 72–74 (describing “connection properties” including “transaction ordering” and “bounds on connection throughput, latency, and jitter,” and “transaction

completion”). Drake likewise describes that its channels have properties that include “throughput” and “end-to-end delay” (*i.e.*, latency). Ex-1006 (Drake) at 5:30–50.

6. Claim 6

a. Claim 6[pre]: “Integrated circuit according to claim 2”

192. Claim 6 depends from claim 2, which is obvious for the reasons that I explained above.

b. Claim 6[a]: “wherein said first module (M; I) is adapted to issue a request (REQ) for a connection with at least one of said second modules to said communication manager (CM)”

193. In my opinion, Goossens2003 in view of Drake renders this limitation obvious. For example, Goossens2003 describes that “IPs negotiate with network interfaces to obtain connections with certain properties.” Ex-1007 (Goossens2003) at 77. Therefore, in my opinion, a POSITA would have understood that Goossens2003’s master IPs (*i.e.*, first module) issue requests for connections to the network interfaces (*i.e.*, communication managers).

194. As I previously explained above, the requested connections are between a “master” (*i.e.*, first module) and “one or more slaves” (*i.e.*, second modules). Ex-1007 (Goossens2003) at 71. Goossens2003 also explains that when IPs make a request for a connection, the desired communication properties and level of service “must be stated explicitly.” Ex-1007 (Goossens2003) at 67; *see also* Ex-1007 (Goossens2003) at 71 (“[c]onnections must be created stating the requested service

level”). Therefore, in my opinion, a POSITA would have understood that the requests issued by the IPs comprise desired connection properties.

195. As I previously described above in Ground I, Drake similarly describes that an application in a source station (*i.e.*, first module) “requests a QoS connection” and that it directs these requests to a “QoS requestor manager” (*i.e.*, communication manager).

- c. **Claim 6[b]: “said communication manager (CM) is adapted to forward said request (REQ) for a connection with communication channels each having a specific set of connection properties to said resource manager (RM)”**

196. In my opinion, Goossens2003 in view of Drake renders this limitation obvious. As I previously discussed above in Ground 1, Drake explains that, having received a request for a connection, the QoS requestor manager (*i.e.*, communication manager) “assemble[s] a request to reserve a QoS connection” and forwards the request to the QoS allocator (*i.e.*, resource manager). Additionally, as I discussed above with respect to claim element 1[e], both Goossens2003 and Drake describe that the requested connections comprise communication channels each having a specific set of connection properties. *See* above Claim Element 1[e].

d. Claim 6[c]: “said resource manager (RM) responding with the availability of the requested connection to said communication manager (CM)”

197. In my opinion, Goossens2003 in view of Drake renders this limitation obvious. As I previously discussed above in Ground I, Drake describes that, when the the QoS allocator receives requests for connections with certain QoS properties, it “examine[s] the resources” of the network to determine whether those requests can be satisfied and sends a “QoS allocation reply” to the QoS requestor manager (*i.e.*, communication manager) which indicates, among other things, whether the requested connection and connection properties are available.

e. Claim 6[d]: “wherein a connection between the first and second module is established based on the available properties of said communication channels of said connection”

198. In my opinion, Goossens2003 in view of Drake renders this limitation obvious. For example, Goossens2003 describes that a request for a connection will be granted and the connection established if the NoC “has sufficient resources available.” Ex-1007 (Goossens2003) at 67. Similarly, as I previously discussed above in Ground I, Drake describes that, if the QoS allocator determines that the requested QoS properties are available, the requested connection between the source and target stations is established and required resources are reserved.

7. Claim 9: “Integrated circuit according to claim 2, further comprising: at least one network interface means (NI), associated to each of said modules, for managing the communication between said modules and said network (N)”

199. Claim 9 depends from claim 2, which is obvious for the reasons that I explained above. Additionally, in my opinion, this limitation is disclosed by Goossens2003. Goossens2003 explains that, in its SOC, “network interfaces” translate between the “transport layer” utilized by the IPs to request services and the “network layer” utilized by the routers of the NOC to transport data. Ex-1007 at 74; *see also* Ex-1007 (Goossens2003) at 75 (describing “network interface to network interface” communications). Goossens2003 also explains that “IPs negotiate with network interfaces to obtain connections with certain properties” and reserve network resources. Ex-1007 (Goossens2003) at 76–77. Therefore, in my opinion,, a POSITA would have understood that these network interfaces manage the communications between the IPs (*i.e.*, modules) and the NOC (*i.e.*, network).

200. The position of these network interfaces between the IP modules and the NOC is shown in Goossens2003’s Figure 4.5, which I have annotated below:.




Figure 4.5. A narrowcast connection.

Ex-1007 at Fig 4.5 (annotated).

8. Claim 10

201. In my opinion, Goossens2003 in view of Drake teaches all the limitations of claim 10, and thus would have rendered that claim obvious as I discuss below.

- a. **Claim 10[pre.1]: “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”**

202. In my opinion, to the extent it is limiting, this part of the preamble is disclosed by Goossens2003. As I previously discussed above with respect to claim elements 1[pre]–1[b] and 1[d], Goossens2003 describes a system on chip (*i.e.*, integrated circuit) in which multiple intellectual property blocks (*i.e.*, modules) communicate with one another over connections via a “network on chip.” *See* above Claim Elements 1[pre]–1[b] and 1[d].

- b. Claim 10[pre.2]: “wherein said connections comprises a set of communication channels each having a set of connection properties, any communication channel being independently configurable”**

203. In my opinion, to the extent it is limiting, this part of the preamble is substantively the same as element 1[f] of claim 1, and is disclosed by Goossens2003 and/or rendered obvious by Goossens2003 in view of Drake for the same reasons that I explained with respect to that claim element. *See* above Claim Element 1[e]. In my opinion, Goossens2003 alone or in view of Drake renders this limitation obvious, including under Patent Owner’s proposed construction of “independently configurable.” Ex-1030, at 13.

- c. Claim 10[pre.3]: “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”**

204. In my opinion, to the extent it is limiting, this part of the preamble is substantively the same as element 1[f] of claim 1, and is disclosed by Goossens2003 for the same reasons that I explained with respect to that claim element. *See* above Claim Element 1[f].

- d. Claim 10[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”**

205. In my opinion, this limitation is substantively the same as element 6[a] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See above Claim Element 6[a].*

- e. Claim 10[b]: “the communication manager forwarding the request to a resource manager”**

206. In my opinion, this limitation is substantively the same as element 6[b] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See above Claim Element 6[b].*

- f. Claim 10[c]: “the resource manager determining whether a target connection with the desired connection properties is available”**

207. In my opinion, this limitation is substantively the same as element 1[g] of claim 1, and is obvious for the same reasons that I explained with respect to that claim element. *See above Claim Element 1[g].*

- g. Claim 10[d]: “the resource manager responding with the availability of the target connection to the communication manager”**

208. In my opinion, this limitation is substantively the same as element 6[c] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See above Claim Element 6[c].*

h. Claim 10[e]: “the target connection between the first and second module being established based on the available properties of said communication channels of said connection”

209. In my opinion, this limitation is substantively the same as element 6[d] of claim 6, and is obvious for the same reasons that I explained with respect to that claim element. *See above Claim Element 6[d].*

9. Claim 11: “The method of claim 10, wherein the network manages traffic utilizing at least one of a switch and a router”

210. Claim 11 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 12 is also rendered obvious by the combination of Goossens2003 and Drake for the reasons that I explained previously with respect to claim 3, which includes an identical limitation. *See above Claim 3.*

10. Claim 12: “The method of claim 10, wherein the processing modules and the network are disposed on a chip”

211. Claim 12 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 12 is also rendered obvious by the combination of Goossens2003 and Drake for the reasons that I explained previously with respect to claim 4, which includes an identical limitation. *See above Claim 4.*

11. Claim 13: The method of claim 10, wherein the connection properties comprise at least one of data transport ordering, flow control, throughput, latency, and lossiness”

212. Claim 13 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 13 is also rendered obvious by the combination of Goossens2003 and Drake for the reasons that I explained previously with respect to claim 5, which includes an identical limitation. *See* above Claim 5. In my opinion, Goossens2003 in view of Drake further discloses the claimed “flow control” property under Patent Owner’s proposed construction (Ex-1019, at 6), as Goossens2003 describes “reserving buffer resources at the receiver” and “using end-to-end flow control to constrain the sender to never send more than the available space.” Ex-1007, at 69.

12. Claim 14: “The method of claim 10, further comprising the resource manager determining whether the target connection with the desired connection properties is available based on reading of a centralized or distributed property table comprising properties associated with the network”

213. Claim 14 depends from claim 10, which is obvious for the reasons that I explained above. Additionally, in my opinion, this limitation is rendered obvious by Goossens2003 in view of Drake. As I previously discussed above in Ground I, Drake describes that its QoS allocator (*i.e.*, resource manager) determines whether a requested connection is available by referencing a “QoS allocation data base” (*i.e.*, a centralized property table), which contains information regarding the properties

associated with the components of the network and the resources reserved by previously allocated connections.

13. Claim 15: “The method of claim 10, further comprising the resource manager reserving the target connection after determining whether the target connection with the desired connection properties is available”

214. Claim 15 depends from claim 10, which is obvious for the reasons that I explained above. Additionally, in my opinion, as I previously discussed above in Ground I, Drake describes that, when the QoS allocator (*i.e.*, resource manager) determines that a requested connection is available, the QoS allocator reserves the resources required for the requested connection. Goossens2003 similarly explains that “[c]onnection acceptance may lead to resource reservations in the NOC.” Ex-1007 (Goossens2003) at 71.

14. Claim 16: “The method of claim 10, wherein the connection properties comprise at least one of throughput, latency and jitter”

215. Claim 16 depends from claim 10, which is obvious for the reasons that I explained above. In my opinion, the additional limitation of claim 16 is also rendered obvious by the combination of Goossens2003 and Drake for the reasons I explained with respect to claim 5. *See* above Claim 5.

216. As I explained above, Goossens2003 describes that each of its channels has properties that include at least “throughput guarantees.” Ex-1007 (Goossens2003) at 74; *see also* Ex-1007 (Goossens2003) at 72–74 (describing

“connection properties” including “bounds on connection throughput, latency, and jitter,” and “transaction completion”). Drake likewise describes that its channels have properties that include “throughput” and “end-to-end delay” (*i.e.*, latency). Ex-1006 (Drake) at 5:30–50.

X. SECONDARY CONSIDERATIONS


217. I am unaware of any secondary considerations of non-obviousness that provide evidence that the Challenged Claims of the '449 patent are not obvious. Should Patent Owner assert any secondary considerations of non-obviousness, I reserve the right to respond at that time.

XI. CONCLUSION

I hereby declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct, and that all statements made of my own knowledge are true and that all statements made on information and belief are believed to be true. I understand that willful false statements are punishable by fine or imprisonment or both under Section 1001 of Title 18 of the U.S. Code.

Date: December 18, 2023

Respectfully submitted,

A handwritten signature in black ink, appearing to read "Todd Mowry", written over a horizontal line.

Todd Mowry, Ph.D.