

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner

IPR2024-00503
U.S. Patent No. 7,035,279

**PETITION FOR *INTER PARTES* REVIEW
UNDER 35 U.S.C. § 312 AND 37 C.F.R. § 42.104**

TABLE OF CONTENTS

Petitioner’s Exhibit List	5
I. Introduction.....	7
II. Grounds for standing	7
III. Note.....	7
IV. Summary of the ’279 patent	8
V. Prosecution history	10
VI. Effective priority date of the ’279 patent	11
VII. Level of ordinary skill in the art	12
VIII. Claim construction.....	12
IX. Relief requested and the reasons for the requested relief.....	12
X. Identification of how the claims are unpatentable.....	13
A. Challenged claims	13
B. Statutory grounds for challenges.....	13
C. Ground 1	16
1. Kovvali.....	16
2. Kalman	20
3. Reasons to Combine	23
4. Claim 1	26
5. Claim 2.....	50
6. Claim 3	51

7.	Claim 4.....	53
8.	Claim 5.....	55
9.	Claim 9.....	56
10.	Claim 10.....	57
11.	Claim 11.....	59
12.	Claim 12.....	60
13.	Claim 13.....	60
14.	Claim 14.....	61
15.	Claim 15.....	63
16.	Claim 16.....	65
17.	Claim 17.....	65
18.	Claim 18.....	67
D.	Ground 2.....	67
1.	Masters.....	67
2.	Reasons to Combine	71
3.	Claim 6.....	73
E.	Ground 3.....	75
1.	RFC2892	75
2.	Reasons to Combine	75
3.	Claim 7.....	76
4.	Claim 8.....	78

XI.	Discretionary denial is inappropriate.....	80
A.	No basis for 35 U.S.C. § 325(d) denial.	80
B.	Discretionary denial under <i>Fintiv</i> is not appropriate.	81
1.	Whether evidence exists that a stay may be granted if a proceeding is instituted	81
2.	Proximity of the court’s trial date to the board’s projected statutory deadline for a final written decision	82
3.	Investment in the parallel proceeding.....	83
4.	Overlap between issues raised in the petition and in the parallel proceeding.....	83
5.	Whether the petitioner is a defendant	83
6.	Other circumstances that impact the Board’s exercise of discretion, including the merits.....	84
C.	Discretionary denial under <i>General Plastic</i> is not appropriate.....	85
XII.	Conclusion	85
XIII.	Mandatory notices	86
A.	Real party-in-interest.....	86
B.	Related matters	86
C.	Lead and back-up counsel and service information.....	86
	Certificate of Word Count	88
	Certificate of Service	89

PETITIONER’S EXHIBIT LIST

Ex.1001	U.S. 7,035,279 (“Bruckman”)
Ex.1002	Prosecution History of U.S. 7,035,279
Ex.1003	Declaration of Nader F. Mir, Ph.D. under 37 C.F.R. § 1.68
Ex.1004	<i>Curriculum Vitae</i> of Nader F. Mir, Ph.D.
Ex.1005	U.S. 6,920,113 (“Kovvali”)
Ex.1006	U.S. 6,680,912 (“Kalman”)
Ex.1007	U.S. 5,920,697 (“Masters”)
Ex.1008	U.S. 6,839,767 (“Davies”)
Ex.1009	U.S. 6,424,624 (“Galand”)
Ex.1010	U.S. 6,157,955 (“Narad”)
Ex.1011	Complaint, Portsmouth Network Corp. v. Cisco Sys., Inc., No. 2:23-cv-00441 (E.D. Tex. Sept. 26, 2023)
Ex.1012	Notice of Disclosures, Portsmouth Network Corp. v. Cisco Sys., Inc., No. 2:23-cv-00441 (E.D. Tex. Jan. 8, 2024)
Ex.1013	Order, Portsmouth Network Corp. v. Cisco Sys., Inc., No. 2:23-cv-00441 (E.D. Tex. Dec. 19, 2023)
Ex.1014	Docket Control Order, Portsmouth Network Corp. v. Cisco Sys., Inc., No. 2:23-cv-00441 (E.D. Tex. Dec. 18, 2023)
Ex.1015	U.S. 6,072,773 (“Fichou”)
Ex.1016	U.S. 6,111,673 (“Chang”)
Ex.1017	U.S. 6,161,137 (“Ogdon”)
Ex.1018	Tsiang, D. <i>et al.</i> , Request for Comments 2892 (“RFC2892”), “The Cisco SRP MAC Layer Protocol” (Aug. 2000), <i>available at</i> https://www.rfc-editor.org/info/rfc2892
Ex.1019	Bradner, S., Request for Comments 2026 (“RFC2026”), “The Internet Standards Process – Revision 3” (Oct. 1996), <i>available at</i>

	https://www.rfc-editor.org/info/rfc2026
Ex.1020	U.S. 6,738,819 (“Li”)
Ex.1021	Vidal, K., “Interim Procedure For Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation” (Jun. 21, 2022), <i>available at</i> https://www.uspto.gov/sites/default/files/documents/interim_proc_discretionary_denials_aia_parallel_district_court_litigation_memo_20220621_.pdf

I. INTRODUCTION

Pursuant to 35 U.S.C. §§ 311, 314(a), and 37 C.F.R. § 42.100, Cisco Systems, Inc. (“Petitioner”) respectfully requests that the Board review and cancel as unpatentable under (pre-AIA) 35 U.S.C. §103(a) claims 1-18 (hereinafter, the “Challenged Claims”) of U.S. 7,035,279 (the “’279 patent,” Ex.1001).

The ’279 patent describes “controlling flow allocation in a network that is arranged in a ring topology.” Ex.1001, 2:6-8. As shown below and confirmed in the Declaration of Nadar F. Mir, Ph.D. (Ex.1003) these concepts were well known. For example, Kovvali (Ex.1005) teaches a ring network in which communication may travel on a path in a clockwise or counterclockwise direction, while Kalman (Ex.1006) and Masters (Ex.1007) describe further details about choosing between these path options, and RFC2892 (Ex.1018) specifies known protocols referenced in Kovvali and Kalman. Therefore, the prior art presented in this Petition teach and render obvious the Challenged Claims.

II. GROUNDS FOR STANDING

Petitioner certifies that the ’279 patent is eligible for IPR and that Petitioner is not barred or estopped from requesting IPR challenging the patent claims. 37 C.F.R. § 42.104(a).


III. NOTE

Petitioner cites to exhibits’ original page numbers. **Emphasis** in quoted

material has been added. Claim terms are presented in *italics*.

IV. SUMMARY OF THE '279 PATENT

The '279 patent relates to “flow allocation in a ring topology.” Ex.1001, Title. A network ring is a communication network comprising a plurality of nodes interconnected by links in a ring topology. Ex.1001, 2:19-21, 4:36-38. Each node communicates with every other node over either a counterclockwise path or a clockwise path around the ring. Ex.1001, 4:40-42. Figure 1 of the '279 patent illustrates network 20 that includes nodes 22 interconnected by links 1-6 to form counterclockwise (inner ring 24) and clockwise (outer ring 26) paths:



Ex.1001, Fig. 1 (annotated).

Figure 1 also depicts the communication network may include a dispatcher 28. The dispatcher is responsible for allocating resources (e.g. bandwidth) to nodes 22. Ex.1001, 4:46-55. The dispatcher also determines on which of the rings (i.e. counterclockwise or clockwise) the data flow is to be directed. Ex.1001, 6:29-31. The dispatcher may be separate from the network nodes or may be associated with one of the nodes (e.g. a software process running on the node). Ex.1001, 4:47-51.

When a node receives a request from a client to send data to another node, the node first checks to determine whether it has sufficient resources already

allocated to it to service the request. Ex.1001, 6:11-14. If the allocated resources are inadequate, the node submits a request to the dispatcher for additional resources. Ex.1001, 6:14-17.

If resources are available (both sufficient unallocated resources left in the specific node budget, and sufficient unused resources left on the appropriate segments of at least one of the rings) the dispatcher determines on which of the rings to allocate the resources and allocates one or more additional “quanta” of resources to the node on the selected ring. Ex.1001, 6:29-31, 6:19-33. The node then carries out the data transmission. Ex.1001, 6:33-34.

On the other hand, if the node determines that the resources allocated already to the node are sufficient, the node carries out the requested transmission “without having to communicate with the dispatcher.” Ex.1001, 6:48-51.

When both rings have sufficient resources available on all relevant links, the dispatcher chooses one of the rings on which to direct the data flow. Ex.1001, 8:1-6. “[T]he dispatcher proceeds to choose the ring that has, generally speaking, the lower level of current resource allocation.” Ex.1001, 8:3-5. In one embodiment, the ’279 patent describes that relative resource “usage levels” can be compared for all relevant segments on both rings to select a path for the data flow. Ex.1001, 8:9-11.

V. PROSECUTION HISTORY

The application which led to the ’279 patent was filed on January 9, 2001.

The Office issued a non-final office action on July 29, 2004, rejecting all but two dependent claims as being obvious over several prior art references. *See* Ex.1002, 90-101. In response, the Patent Owner amended certain claims to include subject matter of the dependent claims deemed to recite patentable subject matter, combined other claims along with “additional clarifying language,” and added new claims. *See* Ex.1002, 118-120. The Office then rejected each of these claims as obvious over additional prior art in another non-final office action. *See* Ex.1002, 131-150. Following an applicant initiated interview, Ex.1002, 179-181, Patent Owner withdrew certain claims to overcome a restriction requirement and amended the remaining independent claims to “clarify” that the alleged invention specifies that when resources already allocated to a first node are sufficient, the data flow is directed to a second node “*without requesting additional resources.*” Ex.1002, 182-191. The Office then issued a notice of allowance. Ex.1002, 199-202.

VI. EFFECTIVE PRIORITY DATE OF THE '279 PATENT

The '279 patent issued from an application filed January 9, 2001. The '279 patent does not claim priority to any earlier application, so Petitioner has not undertaken a priority date analysis. Petitioner does not waive any right or opportunity it may have to dispute the priority date of the '279 patent in this or another forum where the issue is relevant.

VII. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill in The Art (“POSITA”) in January 2001 would have had a working knowledge of the network communication technologies that are pertinent to the ’279 patent, including network flow allocation in a ring topology. A POSITA would have had a master’s degree in computer science, computer engineering, or an equivalent, and two years of professional experience relating to network communications and resource allocations. Additional work experience can substitute for specific educational background, and vice versa. Ex.1003, ¶20.

VIII. CLAIM CONSTRUCTION

Claim terms in IPR are construed according to their “ordinary and customary meaning” to those of skill in the art. 37 C.F.R. § 42.100(b). Petitioner submits that, for the purposes of this proceeding and the grounds presented herein, no claim term requires express construction. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017).

IX. RELIEF REQUESTED AND THE REASONS FOR THE REQUESTED RELIEF

Petitioner asks that the Board institute a trial for *inter partes* review and cancel the Challenged Claims in view of the analysis below.

X. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

A. Challenged claims

Petitioner challenges claims 1-18. The Challenged Claims include all claims of the '279 patent. Thus, a finding that the Challenged Claims are unpatentable in this proceeding will resolve the parties' dispute and obviate the need for a trial regarding the '279 patent in the co-pending litigation, substantially reducing the time and expense of that litigation for all parties.

B. Statutory grounds for challenges¹

Grounds	Claims	Basis
#1	1-5, 9-18	§ 103 (Pre-AIA) over Kovvali and Kalman
#2	6	§ 103 (Pre-AIA) over Kovvali, Kalman, and Masters
#3	7-8	§ 103 (Pre-AIA) over Kovvali, Kalman, and RFC 2892

U.S. 6,920,113 (Ex.1005, “**Kovvali**”) was filed March 28, 2000.

U.S. 6,680,912 (Ex.1006, “**Kalman**”) was filed March 3, 2000.

U.S. 5,920,697 (Ex.1007, “**Masters**”) was filed July 11, 1996 and issued

¹ Petitioner relies on the combined teachings of the references and not on a physical incorporation of elements. *See In re Mouttet*, 686 F.3d 1322, 1332 (Fed. Cir. 2012).

July 6, 1999.

D. Tsiang, et al., Request for Comments 2892, “The Cisco SRP MAC Layer Protocol” (Ex.1018, “**RFC2892**”), bears a copyright date of 2000, and is dated August 2000.

Kovvali, Kalman, and Masters are prior art under pre-AIA 35 U.S.C. § 102(e).² Masters is also prior art under pre-AIA § 102(b).

RFC2892 is prior art under pre-AIA § 102(a). RFC2892 is a “Request for Comments” (RFC) document published by the Internet Engineering Task Force (IETF) and qualifies as a prior art “printed publication.” *See, e.g.*, Ex.1019, 6, (“RFCs can be obtained from a number of Internet hosts”), 8 (IETF documents are “readily available to a wide audience”), 26 (IETF participants “shall publicly announce...every activity” relating to the standardization process); Ex.1003, ¶37. The Board has repeatedly found IETF documents, including RFCs, to be “printed publications.” *See, e.g., Apple, Inc. v. VirnetX, Inc.*, IPR2017- 00337, Paper 31, 46-47 (May 30, 2018) (RFCs are “precisely the type of documents whose main purpose is for public disclosure”); *Riot Games, Inc. v. Paltalk Holdings, Inc.*, IPR2018-00130, Paper 11, 30-33 (May 15, 2018) (RFCs are printed publications).

² Even if Patent Owner argues that the ’279 patent is somehow an AIA patent, the references would still qualify as prior art under post-AIA 35 U.S.C. § 102(a).

RFC2892 is a document published by the IETF describing Spatial Reuse Protocol networking standards and improvements, and a POSITA would have reasonably relied upon it. In fact, POSITAs frequently cited these same types of documents before the '279 priority date. Ex.1005, 7:1-4 (citing RFC 1483). The '279 patent acknowledges that RFC2892 is a prior art printed publication. Ex.1001, 1:21-24 (“One solution ... is the Spatial Reuse Protocol (SRP) ... described by Tsiang et al., in Request for Comments (RFC) 2892 of the Internet Engineering Task Force (IETF). This document which is available at www.ietf.org/rfc.html, is incorporated herein by reference.”; “using methods known in the art, as specified by SRP (RFC 2892)”); *see also* Ex.1001, 1:56-59 (citing RFC 2328), 5:53-58 (citing RFC 2005), 6:1-4 (citing RFC 2814)). Ex.1003, ¶37.


Petitioner also cites to additional prior art (e.g. Ex.1008-1010, 1015-1017, 1019-1020) as evidence of the background knowledge of a POSITA and to provide contemporaneous context to support Petitioner’s assertions regarding what a POSITA would have understood from the prior art in the grounds. *See Yeda Research v. Mylan Pharm. Inc.*, 906 F.3d 1031, 1041-1042 (Fed. Cir. 2018) (affirming the use of “supporting evidence relied upon to support the challenge”); 37 C.F.R. § 42.104(b).

C. Ground 1

1. Kovvali

Like the '279 patent, Kovvali describes a network comprising a plurality of “nodes” interconnected by “links” in a “ring” topology. Ex.1005, 8:19-21. Nodes communicate with each other by passing data over dynamic channels around a ring. Ex.1005, 8:39-41. The ring may be of various types including, for example, a UPSR (unidirectional path switched ring) or a BLSR (bi-directional line switched ring). Ex.1005, 8:30-36; Ex.1003, ¶39.

Kovvali’s Figure 4 illustrates a plurality of nodes 420 (labeled A-D) coupled by optical links 422 to form a UPSR, in which all communication between nodes travels in one direction. Ex.1005, 8:30-32.




Ex.1005, Fig. 4 (modified and annotated); Ex.1003, ¶41.

In Figure 4, Kovvali illustrates that node A may be an “arbiter” node that acts as a centralized bandwidth manager for bandwidth requests. Ex.1005, 9:1-6, 7:32-33, 4:47-50. Like the dispatcher in the ’279 patent, the arbiter node is responsible for assigning resources (such as “dynamic channels”) to client nodes and adjusting their allocated rates. Ex.1005, 9:8-11. After the network ring is provisioned, the arbiter node provides “initial data rate allocations for each of the dynamical channels” it assigned. Ex.1005, 12:5-6, Ex.1003, ¶42.

The dynamic channel allocations are in “units of columns” of a dynamic

frame. Ex.1005, 12:6-7. In one example, Kovvali describes the dynamic frame as including three sections: (A) a three-column overhead section 520 for transmitting dynamic bandwidth management information between nodes 420 and the arbiter node, Ex.1005, 9:49-50, (B) a TDM (time division multiplexed) section 530 for statically configured, fixed-rate channels, Ex.1005, 9:57-64, and (C) an elastic section 540 for dynamic channels, Ex.1005, 9:65-10:3, as illustrated in Figure 5A:



Ex.1005, Fig. 5A (annotated); Ex.1003, ¶43.

The dynamic channels in the elastic section can be dynamically allocated to accommodate “bursty data” that may exhibit time-varying rates. Ex.1005, 9:65-10:1, 8:49-53. A client node may request to modify the capacity of assigned channels using control channels in the dynamic frame overhead. Ex.1005, 10:14-18. For example, a client node may send a request for additional bandwidth to the

arbiter node by setting “a request section (REQ)” in the dynamic frame overhead for a bursty channel. Ex.1005, 10:35-38. The request may be for zero or more columns of bandwidth in the dynamic frame. Ex.1005, 10:39-54, Ex.1003, ¶44.

“The arbiter node then determines how to increase or decrease allocations to one or more dynamic channels to accommodate the request.” Ex.1005, 12:35-37. The arbiter node indicates the changes in allocations by setting a “grant section (GNT)” in the dynamic frame overhead for the corresponding bursty channel, indicating zero or more columns of allocated bandwidth. Ex.1005, 10:37-38, 10:55-60; Ex.1003, ¶45.

2. Kalman

Kalman relates to directing traffic in a ring network topology. Kalman describes a technique for selecting “the preferred output direction for traffic from the source to any destination” on a bi-directional ring network. Ex.1006, 8:5-7. Like the '279 patent, “[t]he disclosed network includes two rings, wherein a first ring transmits data in a ‘clockwise’ direction (or direction 0), and the other ring transmits data in a ‘counterclockwise’ direction (or direction 1 opposite to direction 0).” Ex.1006, 2:61-65. Although there are two paths between any two nodes (i.e. clockwise and counterclockwise), “[a] given packet/flow between two nodes is transmitted in only a single direction around the network (even when there is a span fault) and is removed from the ring by the destination node.” Ex.1006,

4:40-43, Ex.1003, ¶46.

Figure 3 illustrates an example “where node 0 transmits information to node 5 in only the [counterclockwise] direction.... A transmission from node 5 to node 0 would only go through nodes 6 and 7 in the opposite direction. This allows for optimized ring capacity utilization since no capacity is set aside for protection.” Ex.1006, 4:43-48. At node 5, the transmitted information is removed from the ring. The portion of bandwidth used for the transmission from node 0 to node 5 is thus available for other transmissions on the counterclockwise path from node 5 back to node 0. Kalman describes this as “spatial reuse of capacity.” Ex.1006, 4:58-67, Ex.1003, ¶47.




FIG. 3

Ex.1006, Fig. 3 (annotated); Ex.1003, ¶47

Typically, the “least-cost physical route” is selected as the preferred direction for a given data flow. Ex.1006, 4:49-50. Kalman describes that “[f]or each destination node j , a cost is assigned to each output direction (0 and 1)” from the source node on the ring. Ex.1006, 7:50-51. A preferred direction is selected “based on the direction with the lowest cost.” Ex.1006, 7:51-53; Ex.1003, ¶48.

The computation of cost from each source node to each destination node “may take into account a variety of factors, including the number of hops from

source to destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and the level of traffic congestion in that direction.” Ex.1006, 7:66-8:5. “Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly.” Ex.1006, 8:5-7; Ex.1003, ¶49.

3. Reasons to Combine

A POSITA would have combined concepts from Kovvali and Kalman because they provide complementary ideas for the efficient transmission of data in a ring network. *See, e.g.*, Ex.1005, Abstract; Ex.1006, 4:20-23; Ex.1003, ¶50.

A POSITA would have been motivated to combine Kovvali with Kalman, and doing so would have been within a POSITA’s skillset. The combination would have been obvious because it represents merely the use of Kalman’s known technique for selecting a preferred path on a ring network to improve Kovvali’s multi-path ring network for transmitting dynamic frames in the same way. Faced with the choice of transmitting a dynamic frame over one of two paths (e.g. clockwise path or counterclockwise path), a POSITA would look to Kalman for details on how to make the path selection. Ex.1003, ¶51.

The combination would have also been obvious because the application of Kalman’s technique (selecting a preferred path on a ring network) to improve Kovvali’s method (transmitting dynamic frames on a multi-path ring) ready for

improvement would yield predictable results. Here, the improvement provided to Kovvali's multi-path ring network by Kalman's path selection technique is efficient transmission of dynamic frames in a ring network using a lowest-cost path. Ex.1003, ¶52.

Kovvali's multi-path ring network would also be improved because the use of a ring selection technique by a centralized bandwidth manager (i.e. Kovvali's "arbiter" node) would allow path selection that accounts for dynamic frame transmissions among all the network nodes across the entire network. By centralizing path selection decisions at the arbiter, more efficient utilization of available network capacity is achieved. Ex.1003, ¶53.

For example, referring to Kovvali's modified and annotated Figure 4 discussed in Section X.C.1, in response to a first request to transmit a first dynamic channel from Node B to Node D, the lowest cost path calculated for transmission may be the clockwise (blue) path through Node C. Kalman teaches that the lowest cost path is "typically" used for transmission of "a given packet/flow" between nodes, Ex.1006, 4:49-50, 4:40-45, and thus the arbiter would select this clockwise path for transmitting the data flow. Generally, however, the arbiter would expect to receive more than one request for a dynamic channel, such as a second request for a second dynamic channel to be transmitted between Node B and Node C. In such a scenario, when selecting a path for transmission of a particular dynamic channel,

the arbiter would consider other factors such as the priority of data to be transmitted. For example, if the data associated with the first request has low priority and is delay tolerant, the arbiter may choose to send the respective dynamic channel over a more congested path thereby leaving available paths with less congestion for other requests that may have higher priority and are less delay tolerant. *See, e.g.*, Ex.1006, 15:15-19, 16:49-50 (noting how some “lower priority traffic” may be “delayed” to leave bandwidth available to support transmission requirements of higher-priority “protected” traffic). Thus, by maintaining information about the current amount of bandwidth allocation throughout the network (e.g. congestion), the centralized arbiter would make use of network capacity efficiently by selecting available paths to service multiple requests according to the particular characteristics of the data they transmit. Ex.1003, ¶54.

Further, a POSITA would have understood that engineers working on bandwidth allocation schemes would consider a myriad of network configurations to use network resources efficiently. For example, it was well known that, “[t]he final aim of a congestion scheme is to optimize the overall network performances. Schemes that optimize the throughput of one user or one trunk, do not always maximize the overall performance.” Ex.1009, 6:1-5. Thus, a POSITA would have been motivated to implement Kalman’s ring selection technique in Kovvali’s centralized bandwidth manager (the “arbiter” node) to optimize the use of network

resources for all dynamic channels instead of just one. Ex.1003, ¶55.

A POSITA would have had a reasonable expectation of success in combining the teachings of Kovvali and Kalman. Kalman describes that its technique for selecting a preferred path for transmitting data on a network ring may be applied “to any network transmitting frames, cells, or using any other protocol.” Ex.1006, 3:38-41. Kalman also describes that, like Kovvali, the nodes of the network ring are interconnect by fiber optic cables. Ex.1006, 11:37-42. Because Kovvali also describes a ring network transmitting dynamic frames between nodes over an optical network, a POSITA would have understood that the path selection techniques described in Kalman are applicable to Kovvali. A POSITA would also have had an expectation of success from the combination because of the overall similarities between the network architectures and technologies described in Kovvali and Kalman. Ex.1003, ¶56.


4. Claim 1

[1.0] *In a communication network that includes a plurality of nodes interconnected by links, the nodes including at least first and second nodes, a method for routing a data flow, comprising:*

To the extent the preamble is limiting, Kovvali renders it obvious.

Kovvali teaches “a method for communicating between multiple nodes of a communication system in which a series of fixed-length frames pass between the nodes of the system.” Ex.1005, 4:55-58; Ex.1003, ¶58. The nodes may be “IP-

based routers” that route a data flow (e.g. data, voice and/or video “traffic”) in a ring topology. Ex.1005, 4:12-14, 35-40; Ex.1003, ¶59. The “nodes 420” are coupled over “optical links 422,” shown in Figure 4 below. Ex.1005, 8:19-21; Ex.1003, ¶60.



Ex.1005, Fig. 4 (annotated); Ex.1003, ¶60.


Therefore, Kovvali renders the preamble obvious. Ex.1003, ¶¶57, 61.

[1.1] allocating a respective number of quanta of one or more resources in the network to each of the nodes;

First, Kovvali teaches that “one node 420 is assigned to be the arbiter node

at the time the ring is provisioned. The remaining nodes ... are client nodes,”

Ex.1005, 9:1-5; Ex.1003, ¶63, as depicted in Fig. 4:




Ex.1005, Fig. 4 (annotated); Ex.1003, ¶63.

Second, Kovvali describes that during configuration of the network ring, “the arbiter node assigns a number of dynamic channels 436 to couple particular nodes (step 930). This involves the arbiter node notifying each node of the dynamic channel which it terminates, and providing the initial data rate allocations for each of the dynamic channels.” Ex.1005, 12:1-6; Ex.1003, ¶64. The arbiter node also “**assigns dynamic channels to particular nodes**” and “adjust[s] the


allocated rates for the dynamic channel” in response to “requests” from a client node to the arbiter. Ex.1005, 9:6-11; Ex.1003, ¶65.

Third, Kovvali teaches that each dynamic channel is allocated zero or more columns (“*quanta*”) in a frame structure. Ex.1005, 10:3-5; Ex.1003, ¶¶66, 71. In particular, Kovvali describes a “frame” that includes a “Synchronous Payload Envelope (SPE)” for carrying data between nodes. Ex.1005, 2:36-39; Ex.1003, ¶67. Each frame comprises “an array of 90 columns by 9 rows of 8-bit bytes,” Ex.1005, 2:55-58; Ex.1003, ¶67, as depicted in Figure 3 (showing two frames):



Ex.1005, Fig. 3 (annotated); Ex.1003, ¶67.

“Each SPE [Synchronous Payload Envelope] is 87 columns by 9 rows in size,” as depicted in Figure 5A. Ex.1005, 2:66; Ex.1003, ¶68.



Ex.1005, Fig. 5A (annotated); Ex.1003, ¶69.

As shown in Figure 5A, the SPE includes a dynamic frame comprising three sections: (A) a “dynamic frame overhead section 520,” (B) a “TDM section 530,” and (C) an “elastic section 540 [that] carries channels that are dynamically allocated.” Ex.1005, 9:48-66; Ex.1003, ¶70. The dynamic channels in the elastic section are “flexible bandwidth channels” where bandwidth “**allocations are in units of columns** of the dynamic frame.” Ex.1005, 4:43-46, 12:6-7; Ex.1003,

¶¶71-72.

Kovvali's allocation of bandwidth to dynamic channels assigned to nodes is much like the example of allocating "*resources*" in the '279 patent specification, such as "bandwidth" or "any other allowable factor that may affect the ability of a network link to convey a data flow between nodes." Ex.1001, 2:12-18. Thus, Kovvali renders obvious "*allocating a respective number ... of one or more resources in the network to each of the nodes.*" Ex.1003, ¶73.


Furthermore, it would have been obvious to a POSITA that allocating bandwidth in "units of columns" is allocating bandwidth in "*quanta*" because each column is the smallest discrete unit of bandwidth that can be allocated to a dynamic channel. Ex.1003, ¶74. Referring to Figure 5A discussed above, each column of bandwidth corresponds to 9 rows of 8-bit bytes; therefore, each column corresponds to 72 bits. Ex.1005, 2:56-58; Ex.1003, ¶74. Two columns of bandwidth would correspond to 144 bits. Thus, regardless of whether 73 or 144 bits are transmitted, no change in bandwidth allocation is needed because two columns of bandwidth can carry from 73 up to 144 bits of information on the dynamic channel. Ex.1003, ¶74. This teaching is much like the example of allocating resources in "*quanta*" in the '279 patent specification where "[r]esources are allocated in quanta in order to limit the frequency with which allocation and deallocation operations must be performed." Ex.1001, 5:28-30; Ex.1003, ¶75.

Therefore, Kovvali's disclosure of assigning "dynamic channels" to the nodes in the network ring where the resources in the dynamic channels are allocated in "units of columns," renders obvious this limitation. Ex.1003, ¶¶62, 76.

[1.2] *receiving a request at the first node to use a portion of the resources so as to carry the data flow between the first and second nodes over one of a plurality of paths therebetween, each such path comprising a respective sequence of the links;*


First, as discussed in [1.1] "each dynamic channel ... is used for communication between two or more nodes 420" ("*carry data flow between the first and second nodes*"). Ex.1005, 10:3-3; Ex.1003, ¶78.

Second, Kovvali describes that "each node 420 functions as an add/drop multiplexer for the ... dynamic channels. That is, the node ... multiplexes outbound data onto the ring for those channels." Ex.1005, 12:20-24; Ex.1003, ¶79. Each node that terminates (e.g. that adds) a dynamic channel multiplexes outbound data that was received locally for transmission onto the dynamic channel. Ex.1005, 14:33-37; Ex.1003, ¶80. Figure 4 depicts an example of Node B and Node C each receiving data from local dynamic channels 436 for transmission (and correspondingly extracting inbound data from dynamic channels and transmitting it to the local dynamic channels 436) through the use of double-headed arrows (green). Ex.1003, ¶80.



Ex.1005, Fig. 4 (annotated); Ex.1003, ¶80.

Figure 6 depicts another example where a node 420 receives outgoing packet traffic from local dynamic channels 436 for transmission:




Ex.1005, Fig. 6 (annotated); Ex.1003, ¶81.

In Figure 6, a “Tx [transmit] packet processor 634 multiplexes dynamic channels ... that are passed from packet processor 644 onto Tx bus 614, and then to SONET framer 610 where it is multiplexed onto the outbound optical link.” Ex.1005, 14:33-37; Ex.1003, ¶82. Packet processor 644 “provides an interface to [local] dynamic channels 436” and “includes buffers and other mechanisms to implement appropriate queuing, scheduling and policing processes necessary to

handle incoming and outgoing packet traffic.” Ex.1005, 14:27-33; Ex.1003, ¶82.

Accordingly, a POSITA would have understood that the receipt of data from local dynamic channels by a node (e.g. Node B or Node C) for transmission to another node, renders obvious “*receiving a request at the first node to use a portion of the resources*” because it is a request for the node to process (e.g. queue, schedule, police) and multiplex the data onto the portion of network resources allocated by the arbiter (as discussed in [1.1]) to the node for the respective dynamic channels in the dynamic frame. Ex.1003, ¶83.

Third, Kovvali describes that while Figure 4 discussed above depicts “a UPSR (unidirectional path switched ring), in which all communication between nodes travels in one direction,” alternative embodiments include other types of rings such as “a BLSR (bi-directional line switched ring) in which communication between nodes on the ring travels in both clockwise and counterclockwise directions, with two unidirectional optical links joining each pair of neighboring nodes.” Ex.1005, 2:10-16, 8:33-36; Ex.1003, ¶84. Adapting Kovvali’s Figure 4 to use a BLSR (bi-directional line switched ring) as suggested, modified Figure 4 below shows each pair of nodes connected with two unidirectional optical links traveling in opposite directions:



Ex.1005, Fig. 4 (modified and annotated); Ex.1003, ¶85.

As shown in modified Figure 4, above, the resulting network includes two paths for transmitting data between nodes: a first path with a sequence of links in a clockwise direction (orange) and a second path with a sequence of links in a counterclockwise direction (blue). Ex.1003, ¶86.

Therefore, Kovvali's disclosure of a node receiving, processing, and multiplexing data from local dynamic channels onto a ring network that is connected to other nodes with two unidirectional optical links traveling in opposite directions, renders obvious this limitation. Ex.1003, ¶¶77, 87.

[1.3] and in response to the request: if the resources already allocated to the first node on one of the paths are sufficient, directing the data flow from the first node to the second node over the one of the paths without requesting additional resources;

First, as discussed in [1.1], Kovvali describes that an “arbiter node assigns a number of dynamic channels 436” to particular nodes in the network with “initial data rate allocations” during a configuration step. Ex.1005, 12:1-6; Ex.1003, ¶89. The arbiter node may also assign dynamic channels and “adjust[] the allocated rates for the dynamic channels” upon request from respective client nodes. Ex.1005, 9:6-11; Ex.1003, ¶90.

Second, as discussed in [1.2], Kovvali teaches that each node that terminates (e.g. adds or drops) a dynamic channel multiplexes outbound data that was received locally for transmission onto the dynamic channel. Ex.1005, 14:33-37; Ex.1003, ¶91. In particular, a node adding data to a dynamic channel “determines the portion of a frame that is allocated,” “adds data for transmission...into the portion of the...frame that is allocated,” and “transmits the...frame onto the network.” Ex.1005, 5:66-6:12; Ex.1003, ¶92.

A POSITA would have understood that adding (or multiplexing) data for transmission into the portion of a frame that is already allocated to a dynamic channel, renders obvious “*if the resources already allocated to the first node on one of the paths is sufficient, directing the data flow to the second node . . . without*

requesting additional resources” because the data is included in only the portion of the frame that is already allocated to the dynamic channel without requesting additional resources, such as bandwidth. As discussed further at [1.4] below, Kovvali explains that a node may take action (e.g. request additional bandwidth) if the data to be transmitted onto a dynamic channel exceeds the allocated bandwidth; it would have been obvious to a POSITA, therefore, that no action would be needed or taken (e.g. the node would not request additional bandwidth) when the existing allocation is sufficient (e.g., the data to be transmitted fits within the number of columns in the transmission frame already allocated, as discussed at [1.1]). Ex.1003, ¶93. As discussed at [1.3], Kovvali’s arbiter makes “initial data rate allocations” before any node issues a bandwidth request. Ex.1005, 12:1-6. Nodes then use these allocations to coordinate and control how they multiplex their data into their assigned dynamic channels for transmission. Ex.1005, 12:15-27. Thus, it would have been obvious to a POSITA for Kovvali’s nodes to have “*resources already allocated*” and available for use without requiring a request for additional resources. Ex.1003, ¶93.

Third, as discussed in [1.2] the BSLR network taught by Kovvali includes two paths for transmitting between nodes: a clockwise path and a counterclockwise path. Further, as discussed in [1.1] dynamic channels are used for communication between two or more nodes. Thus, a POSITA would understand that Kovvali’s

disclosure of transmitting a frame “onto the network” after multiplexing data for a dynamic channel into a frame renders obvious “*directing the data flow from the first node to the second node over the one of the paths without requesting additional resources,*” because the frame transmitted on the network would use either the clockwise or the counterclockwise path. Ex.1003, ¶94.

Therefore, Kovvali’s disclosure of a node receiving outbound data that was received locally for transmission onto the dynamic channel, multiplexing the outbound data for transmission into the portion of the frame that is already allocated to the dynamic channel without requesting additional resources, and transmitting the frame onto a BSLR network, renders obvious this limitation. Ex.1003, ¶¶88, 95.

[1.4] *and if the resources already allocated to the first node are insufficient, performing the following steps:*

Kovvali explains that data transmitted on the dynamic channels may be “bursty” in nature, and thus a node may take additional steps when it has a large burst of data for transmission. Ex.1005, 4:30-34. It would have been obvious to a POSITA that for such bursty traffic, the data to be transmitted would occasionally exceed the bandwidth allocated for it (and thus, “*the resources already allocated to the first node are insufficient*”). Ex.1003, ¶96. As explained further below, the prior art renders obvious performing the recited steps. *See* [1.5], [1.6], [1.7] and


[1.8] *infra.* Ex.1003, ¶96.

[1.5] *requesting an increase in an allocation to the first node of the requested resources;*

First, Kovvali teaches data traffic may “exhibit bursty or time-varying rates.” Ex.1005, 4:32-34, 8:51-53; Ex.1003, ¶98.

Second, “to accommodate time-varying data rate requirements of bursty communication between nodes,” Kovvali describes that “the data rate assigned to any particular dynamic channel 436 can be dynamically changed.” Ex.1005, 8:49-53; Ex.1003, ¶99.

Third, Kovvali describes that a client node may request an increase in bandwidth for a bursty channel by setting a “request section (REQ) 526” in a control channel 524 sent to the arbiter. Ex.1005, 10:35-37. The arbiter responds to the request by setting a “grant section (GNT) 527” to grant bandwidth to the requested channel. Ex.1005, 10:35-38; Ex.1003, ¶100. Figure 5A depicts an embodiment of a dynamic frame that includes the “request section (REQ) 526” and the “grant section (GNT) 527” as part of a “dynamic frame overhead section 520,” Ex.1005, 9:48-52:



Ex.1005, Fig. 5A (annotated), Ex.1003, ¶101.


Thus, if the bandwidth of a dynamic channel already assigned to a node becomes insufficient because the data rate requirement on the dynamic channel has increased (e.g. the node has a burst of data to be transmitted), it would have been

obvious for the client node to send a request to the arbiter node for additional bandwidth. Ex.1003, ¶102.

Therefore, Kovvali's description of a client node sending a request to an arbiter node for additional bandwidth to accommodate a dynamic channel with an increased data rate requirement, renders obvious "*if the resources already allocated to the first node are insufficient . . . requesting an increase in allocation to the first node of the requested resources.*" Ex.1003, ¶¶97, 103.

[1.6] *determining, for the links comprised by each of the paths, respective levels of use of the requested resources due to communications in progress over the network;*

First, as discussed in [1.2], Kovvali's teaching of a bidirectional ring network provides two paths for transmitting data between nodes, illustrated in modified Figure 4 below. There is a first path in a clockwise direction (blue) and a second path in a counterclockwise direction (orange). It would have been obvious to a POSITA that when the arbiter receives a request to increase the bandwidth allocated for communications from a first node to a second node, the arbiter would have two path options to choose from for potentially accommodating that bandwidth increase request. Ex.1003, ¶105.



Ex.1005, Fig. 4 (modified and annotated); Ex.1003, ¶105.

Second, Kalman discloses a technique for selecting a “preferred direction for traffic” between two nodes, “based on the direction with the lowest cost.” Ex.1006, 7:50-53; Ex.1003, ¶106. Like Kovvali, Kalman discloses a network that “includes two rings, wherein a first ring transmits data in a ‘clockwise’ direction (or direction 0), and the other ring transmits data in a ‘counterclockwise’ direction (or direction 1 opposite of direction 0).” Ex.1006, 2:61-63. “[D]ata between nodes flows on the ring that provides the lowest-cost path to the destination node,” Ex.1006, 3:2-3, “rather than in both directions,” Ex.1006, 4:24-26, Ex.1003, ¶107.

“A given packet/flow between two nodes is transmitted in only a single direction around the network ... as is shown in FIG. 3 where node 0 transmits information to node 5 in only the [counterclockwise] direction A transmission from node 5 to node 0 would only go through nodes 6 and 7 in the opposite direction.” Ex.1006, 4:40-48; Ex.1003, ¶¶108-109.




FIG. 3

Ex.1006, Fig. 3 (annotated); Ex.1003, ¶108.

The computation of cost on a direction from a source node to a destination node “may take into account a variety of factors,” including “the level of traffic

congestion in that direction.” Ex.1006, 7:66-8:5; Ex.1003 ¶¶110-111. A POSITA would have been familiar that “traffic congestion” was a common way to refer to the level of use of link bandwidth and/or other resources in a network by the traffic being sent over the network. Ex.1003, ¶112; *see also* Ex.1008, 8:34-36 (“A congestion condition occurs when the sources 12, 13, attempt to exceed the finite bandwidth of the resources along the paths 10, 12”); Ex.1009, 1:36-39 (“Congestion is generally defined as a condition during which network performance is degraded due to saturation of network resources such as communication links, processor cycles, memory buffers, etc.”). Thus, a POSITA would have found it obvious for the “level of traffic congestion” disclosed by Kalman to include the level of use of link bandwidth and/or other resources by the dynamic frame resources described in Kovvali. Ex.1003, ¶112.

Third, Kalman discloses that a “table” may be maintained indicating the cost of transmitting data from a source node to a destination node in both output directions (i.e. “direction 0 cost, and direction 1 cost”) and the preferred direction. Ex.1006, 7:64-65; Ex.1003, ¶¶113-114. Table 1 is an example of such table showing “[t]he computation of cost on a direction from node 0 (assuming node 0 as the source) to node j”:

TABLE 1

Preferred direction table at node 0			
Destination Node	Direction 0 cost	Direction 1 cost	Preferred Direction
1	1	7	0
2	2	6	0
3	3	5	0
4	4	4	0
5	5	3	1
6	6	2	1
7	7	1	1

Ex.1006, T.1, 8:21-33; Ex.1003, ¶114.

“Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly.” Ex.1006, 8:5-7. “The preferred direction is that with the lower cost to reach destination node j.” Ex.1006, 8:34-35; Ex.1003, ¶115.

As described in more detail in Section X.C.3, a POSITA would have been motivated to combine the teachings of Kovvali and Kalman, and doing so would have been within the POSITA’s skillset. A POSITA would have appreciated that by having Kovvali’s arbiter direct data transmission toward the least-occupied or less congested portion of the ring network (e.g., to the clockwise or counterclockwise portion that has more unallocated bandwidth), more efficient utilization of the available network capacity would be achieved. Alternatively, a POSITA would have appreciated that having Kovvali’s arbiter direct data

transmission toward the most-occupied or more congested portion of the ring network (e.g., to the clockwise or counterclockwise portion that has more allocated bandwidth) would leave a larger contiguous block of unallocated space on the other portion of the network (e.g., in the opposite direction), thus making it possible to accommodate a larger bandwidth request that may arrive in the future. While either approach would have its respective benefits, consistently applying either strategy would depend upon Kovvali's arbiter node maintaining information about the current amount of bandwidth allocations throughout the network (*"determining, for the links comprised by each of the paths, respective levels of use of the requested resources"*). Ex.1003, ¶116.

Therefore, the combination of Kovvali and Kalman render this limitation obvious. Ex.1003, ¶¶104, 117.

[1.7] selecting which of the paths is to carry the data flow responsive to the determined levels of use of the requested resources on the links comprised in each of the paths;

See analysis at [1.6]. The combination of Kovvali and Kalman teach selecting a lowest-cost path where the cost accounts for the "level of traffic congestion" on the links of the path. See Ex.1006, 7:66-8:5. Therefore, the combination of Kovvali and Kalman render this limitation obvious. Ex.1003, ¶¶118-119.

[1.8] and increasing the allocation of the resources to the first node

on the selected path by a predetermined quantum.

As discussed in [1.5], Kovvali teaches that a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37. In particular, the client node sets a request section (REQ) within the dynamic frame overhead specifying the amount of requested bandwidth. Ex.1005, 10:35-37. The arbiter responds by setting a grant section (GNT) within the dynamic frame specifying the amount of bandwidth granted. Ex.1005, 10:37-38; Ex.1003, ¶121.

In one embodiment, Kalman describes that the values in REQ and GNT may be “encoded” according to a predetermined amount of bandwidth (e.g., a specified number of columns in the dynamic frame structure) to reduce “the number of bits needed for the control channel” as shown in the table below:

REQ	Requested Bandwidth
000	No bandwidth required
001	Deallocate 1 column
010	Deallocate 3 columns
011	Deallocate 6 columns
100	Allocate 1 additional column
101	Allocate 3 additional columns
110	Allocate 6 additional columns
111	Allocate 10 additional columns
GNT	Granted Bandwidth
00000	No Bandwidth allocated
00001	Numerical value of the number
-11110	of columns allocated
11111	All available columns allocated to the channel

Ex.1005, 10:46-61; Ex.1003, ¶122.

A POSITA would have understood that when the requested bandwidth is granted by the arbiter, the bandwidth (e.g., the number of columns) allocated to the client node would be increased (or decreased) by the requested amount (“*a predetermined quantum*”). For example, if a particular dynamic channel was already allocated 5 columns when a client node sends a request for 1 additional column of bandwidth (i.e., a REQ value of 100), then the value of GNT returned by the arbiter would be “00110” (corresponding to a total allocation of 6 columns). The resources allocated to the client node for sending on the dynamic channel would have been increased by 1 column of additional bandwidth. Because the increased allocation was equal to the amount of increase earlier requested by the node, the allocation would have been increased “*by a predetermined quantum*” as

claimed. It would have been obvious to a POSITA for Kovvali's arbiter to respond to an allocation request by adjusting the allocation by the amount requested when the network had sufficient capacity to accommodate the request. Again, in such instances, the allocation adjustment would be "*by a predetermined quantum*" (e.g., the requested amount) as claimed. Ex.1003, ¶123

Kovvali teaches that all bandwidth requests and grants are made in terms of a number of dynamic frame columns, and thus, the allocations are for predetermined, quantized amounts of bandwidth ("*by a predetermined quantum*"). That is, Kovvali's signaling mechanism for bandwidth requests and grants does not allow for allocations to be made on a per-bit or per-byte basis; instead, as discussed at [1.1], requests and allocations are made in 72-bit increments (corresponding to the number of bits in one dynamic frame column). Ex.1003, ¶123.

Therefore, Kovvali renders this limitation obvious. Ex.1003, ¶¶120, 124.

5. Claim 2

[2.0] *A method according to claim 1, wherein the network comprises an Internet Protocol (IP) network.*

Kovvali discloses one approach to passing data between nodes "is called Dynamic Packet Transfer (DPT). This approach uses **IP** as the basic multiplexing technology ... [and] assumes replacement of the existing transport infrastructure at each SONET node with **IP**-based routers. DPT defines a layer 2 protocol called

Spatial Reuse Protocol (SRP), which is used to transport **IP** datagrams between nodes in a ring.” Ex.1005, 4:9-17; Ex.1003, ¶¶126-127.

It would have been obvious to a POSITA for Kovvali’s ring network to employ Internet Protocol (IP), for example by using the Dynamic Packet Transfer capability discussed in Kovvali. A POSITA would have understood that a ring network that uses the IP networking protocol renders obvious an “*Internet Protocol (IP) network*.” Ex.1003, ¶128.

Employing Internet Protocol (IP) for data transmissions on Kovvali’s ring network also would have been obvious because Kovvali specifically focuses on handling “bursty” traffic, and Internet Protocol (IP) was well-known as a “bursty” data communications protocol. Ex.1005, 8:47-53; Ex.1003, ¶129; Ex.1015, 12:50-51 (“inherently bursty traffic (such as IP—Internet Protocol—traffic)”); Ex.1016, 1:35-36 (“packet-switched (that is, bursty) IP traffic”); Ex.1017, 2:16-19 (describing “the typically ‘bursty’ nature of transmissions between nodes of” a “TCP/IP network”).


Therefore, Kovvali renders this limitation obvious. Ex.1003, ¶¶125, 130.

6. Claim 3

[3.0] *A method according to claim 1, wherein the nodes are interconnected in a ring, and wherein the plurality of paths comprises a first path traversing the ring in one direction and a second path traversing the ring in the opposite direction, and wherein selecting which of the paths is to carry the data flow*

comprises selecting the first or the second path.

First, as discussed in [1.2], Kovvali teaches a BLSR network where each pair of nodes are connected with two unidirectional optical links traveling in opposite directions. As shown in modified Figure 4, below, the resulting BLSR network includes a clockwise path (orange) and a counterclockwise path (blue). The two paths are in “*opposite direction[s]*.” Ex.1003, ¶132.



Ex.1005, Fig. 4 (modified and annotated); Ex.1003, ¶132.

Thus, Kovvali renders obvious “*wherein the nodes are interconnected in a*

ring, and wherein the plurality of paths comprises a first path traversing the ring in one direction and a second path traversing the ring in the opposite direction.”

Ex.1003, ¶133.

Second, as discussed in [1.6], Kalman discloses the first path or the second path is selected as the preferred direction “based on the direction with the lowest cost.” Ex.1006, 7:51-53, 8:34-36; Ex.1003, ¶¶134-135.

Therefore, Kovvali and Kalman render this limitation obvious. Ex.1003, ¶¶131, 136.

7. Claim 4

[4.0] *A method according to claim 3, wherein selecting the first or the second path comprises conveying the data flow over an inner or outer data link ring within the network provided by a Spatial Reuse Protocol (SRP).*

As discussed in [2.0], Kovvali discloses one approach to passing data between nodes in a ring “is called Dynamic Packet Transfer (DPT).” Ex.1005, 4:9-10. “DPT defines a layer 2 protocol called **Spatial Reuse Protocol (SRP)**, which is used to transport IP datagrams between nodes in a ring. SRP allows reuse of ring bandwidth across different sections of a ring of nodes.” Ex.1005, 4:15-18; Ex.1003, ¶¶138-139. It would have been obvious to a POSITA for Kovvali’s ring network to transport IP datagrams using a Spatial Reuse Protocol as Kovvali describes. Ex.1003, ¶140.

Kalman also teaches that “increased ring capacity” is enabled “through **spatial reuse** of capacity,” where “removal of packets from the ring by the destination node ensures that traffic does not use more capacity than is necessary to deliver it to the destination node.” Ex.1006, 4:58-61; Ex.1003, ¶141. A POSITA would thus understand that a ring network that enables increased ring capacity through “spatial reuse” when passing data between nodes renders obvious “conveying data flow . . . within the network provided by a *Spatial Reuse Protocol (SRP)*.” Ex.1003, ¶142.

A POSITA would have had an expectation of success in implementing a *Spatial Reuse Protocol (SRP)* in the context of the Kovvali-Kalman combination. Notably, such protocols had already been developed and defined by standards bodies such as the Internet Engineering Task Force, which had issued RFC2892 entitled “The Cisco SRP MAC Layer Protocol.” Ex.1018, 1. RFC2892 defines the “‘Spatial Reuse Protocol’ (SRP) for use with ring based” networks like those in Kovvali and Kalman. Ex.1003, ¶143.

It also would have been obvious to a POSITA to consider Kovvali’s clockwise and counterclockwise ring paths to correspond to “*inner*” and “*outer*” rings, or vice versa. The terms “inner” and “outer” ring were simply used by network engineers to give names to the two separate paths around a bidirectional ring network; the terms were used somewhat like “clockwise” and

“counterclockwise” to distinguish between the two paths. Ex.1003, ¶144. As the ’279 patent confirms, labeling the paths around a ring network as “inner” and “outer” is arbitrary. Ex.1001, 1:27-28; *see also id.* 2:23-25 (“a clockwise or a counterclockwise direction, or in other words,... over an inner or an outer ring”).

Therefore, Kovvali and Kalman each render this limitation obvious.
Ex.1003, ¶¶137, 145.

8. Claim 5

[5.0] *A method according to claim 1, wherein the one or more resources comprise a link bandwidth.*

First, as discussed in [1.1], Kovvali teaches that the arbiter node assigns a number of dynamic channels to “couple” particular nodes. Ex.1005, 12:1-3. Each dynamic channel can be assigned a bandwidth that “is varied dynamically for data traffic by a centralized bandwidth manager which serves as an arbiter for requests for bandwidth.” Ex.1005, 4:46-50; Ex.1003, ¶147.

Second, as discussed in [1.0], Kovvali teaches that the “nodes 420” are “coupled over optical links 422,” Ex.1005, 8:19-21, Fig. 4; Ex.1003, ¶148.

It would have been obvious to a POSITA that allocating bandwidth to a dynamic channel between two nodes effectively allocates that same bandwidth to the dynamic channel on the corresponding links between the two nodes. Ex.1003, ¶149. This is because—as Kalman confirms—the dynamic channel that couples the

nodes also occupies the allocated portion of bandwidth on the link (or links) along the path from the transmitting node to the receiving node. Ex.1003, ¶149; *see* Ex.1006, 4:40-67, Fig.3 (discussing how, with spatial reuse, data transmissions only occupy bandwidth on the links between the transmitter and receiver, leaving the corresponding bandwidth on other links of the ring available for other transmissions).

Therefore, Kovvali renders this limitation obvious. Ex.1003, ¶¶146, 150.

9. Claim 9

[9.0] *A method according to claim 1, wherein selecting which of the paths is to carry the data comprises verifying that a sufficient amount of the requested resources is available to carry the data flow on every one of the links comprised in the selected path.*

As discussed in [7.0] and [8.0], a POSITA would have found it obvious that when the resources requested for transmitting a dynamic channel exceeds the available capacity on any link of a path, the arbiter would either deny the request (e.g. GNT 00000) or return a lower value up to “[a]ll available columns” (e.g. GNT 11111) for the channel in the frame so that the maximum number of columns that can be allocated to the dynamic channel are not exceeded on any of the links in the path (“*verifying that a sufficient amount of the requested resources is available to carry the data flow on every one of the links comprised in the selected path.*”). In other words, Kovvali and Kalman render obvious for the arbiter to deny

a request for bandwidth that cannot be accommodated. In order to deny such requests, it would have been obvious to a POSITA for the arbiter to check whether there is sufficient bandwidth from the source to the destination of the data flow, including checking for sufficient bandwidth on each link that the data flow will traverse. Ex.1003, ¶152. As Kalman teaches, spatial reuse protocol allows a single portion of a dynamic frame to be used for different purposes (different transmissions) on different links around the network ring. *See* Ex.1006, 4:48-67. Thus, it would have been obvious for Kovvali's arbiter node to maintain the bandwidth allocations it has made for each link of the network, and for the arbiter node to confirm when making a new bandwidth allocation that the relevant links have sufficient capacity. Ex.1003, ¶152.

Thus, the combination of Kovvali and Kalman render this limitation obvious. Ex.1003, ¶¶151, 153.

10. Claim 10

[10.0] *A method according to claim 1, wherein selecting which of the paths is to carry the data flow comprises selecting the one of the paths having the lowest level of a predetermined measure of use of the requested resources.*

As discussed in [1.6] and [1.7], the combination of Kovvali and Kalman render obvious a dispatcher that selects the lowest-cost path that accounts for a variety of factors including the “level of traffic congestion” (“*predetermined*

measure of use of the requested resources”) on the links of the path. It would have been obvious to a POSITA that the level of traffic congestion is a “*predetermined*” measure of use because the level of traffic congestion would need to be measured before it would be used in the cost calculation. Ex.1003, ¶155.

A POSITA would have also found it obvious that where the value of all the factors in the lowest cost path calculation are the same except for the “level of traffic congestion,” the path with the lowest “level of traffic congestion” would be selected because it has the lowest cost. Ex.1003, ¶156.

Kalman also provides examples where cost is based on just one factor such as “packet loss rate,” “average packet queuing delay,” or “variation of packet queuing delay.” Ex.1006, 16:36-44. A POSITA would have understood that each of these factors are examples corresponds to a measurement of resource use. For example, packets were commonly “dropped” or “delayed” in a network by a switch or router that was heavily loaded (was experiencing traffic congestion). Ex.1009, 5:65-55; Ex.1008, 3:53-57. As a network became busier, more packets would be dropped or delayed. Thus, Kalman’s measures of “packet loss rate” and “queueing delay” would have been measurements of network usage. Ex.1003, ¶¶157-158. Accordingly, a POSITA would have found it obvious that the calculation of cost could be based on the level of traffic congestion (or another measure of resource use) alone. Ex.1003, ¶158.

Therefore, the combination of Kovvali and Kalman render this limitation obvious. Ex.1003, ¶¶154, 159.

11. Claim 11

[11.0] *A method according to claim 1, wherein receiving the request comprises choosing a dispatcher within the network to manage allocation of the resources, wherein the dispatcher receives and processes the request.*

First, as discussed in [1.1], Kovvali describes that when the network ring is provisioned, “one node 420 is assigned to be the arbiter node,” and the remaining nodes are “client nodes.” Ex.1005, 9:1-5. Accordingly, a POSITA would have found it obvious that the step of “receiving the request” includes assigning one of the nodes in the ring as an “arbiter” node. Ex.1003, ¶¶161-162.

Second, Kovvali describes that the arbiter node (“dispatcher”) is “a centralized bandwidth manager which serves as an arbiter for requests for bandwidth.” Ex.1005, 4:47-50. “[F]or each dynamic channel, the arbiter maintains a list of all terminal nodes for that channel, as well as the current allocation for that channel.” Ex.1005, 13:10-12. Ex.1003, ¶¶163-164.

Third, as discussed in [1.1] and [1.5], Kovvali describes that the arbiter node receives a request (REQ) from “a client node 420 (FIG. 4) to request additional bandwidth for the corresponding bursty channel from the arbiter. GNT 527 is set by the arbiter node to grant bandwidth to the requested channel.”

Ex.1005, 10:33-38; Ex.1003, ¶¶165-166.

Therefore, Kovvali's disclosure of choosing an arbiter that manages network bandwidth, receives requests for dynamic bandwidth, and responds with bandwidth allocation grants, renders obvious this limitation. Ex.1003, ¶¶160, 167.

12. Claim 12

[12.0] A method according to claim 11, wherein choosing the dispatcher comprises choosing one of the nodes to act as the dispatcher.

As discussed in [1.1] and [11.0], Kovvali describes that when the network ring is provisioned, "one node 420 is assigned to be the arbiter node," and the remaining nodes are client nodes. Ex.1005, 9:1-5; Ex.1003, ¶169.

Therefore, Kovvali's disclosure of choosing one of the nodes in the ring to be the arbiter node ("*dispatcher*"), renders obvious this limitation. Ex.1003, ¶¶168, 170.

13. Claim 13

[13.0] A method according to claim 11, wherein requesting the increase in the allocation comprises requesting and receiving the allocation from the dispatcher.

As discussed in [1.5], Kovvali describes that the arbiter node receives a request (REQ) from "a client node 420 (FIG. 4) to request additional bandwidth for the corresponding bursty channel **from the arbiter.**" Ex.1005, 10:35-37 (emphasis added). A grant section "GNT 527 is set **by the arbiter** node to grant bandwidth to

the requested channel.” Ex.1005, 10:37-38 (emphasis added); Ex.1003, ¶172.

Therefore, Kovvali renders this limitation obvious. Ex.1003, ¶¶171, 173.

14. Claim 14

[14.0] *A method according to claim 1, and comprising, upon release of the requested resources: comparing an amount of the resources that have been released to a predetermined threshold; and if the amount is greater than the predetermined threshold, deallocating the predetermined quantum from the first node.*

First, Kovvali describes that “a particular channel may in fact use that channel to pass data for a set of sub-channels.” Ex.1005, 13:1-2. Each sub-channels carries a separate “stream” of data between terminating nodes “already assigned to a dynamic channel.” Ex.1005, 13:41-46; Ex.1003, ¶¶175-176. A POSITA would have found it obvious that due to the “bursty” nature of data traffic (discussed in [1.5]), an existing sub-channel may no longer be needed (e.g. a burst of data packets to transmit on the subchannel have subsided) and may be released. Ex.1005, 11:10-16 (“In operation, from time to time, nodes may initiate a bandwidth allocation step in which ... an existing channel is released if it is no longer needed.”); Ex.1003, ¶177.

Second, as discussed in [1.5], Kovvali describes that to accommodate “bursty” traffic between nodes, a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37, 10:46-55. It would have been obvious to a POSITA that upon release of a sub-

channel, the corresponding bandwidth requirement of the dynamic channel would be reduced by the resources allocated to the subchannel. In this scenario a client node would determine whether to request a decrease in the dynamic channel allocation. Doing so would reduce the resources that would be unnecessarily allocated for a sub-channel that no longer needs them. Ex.1003, ¶¶178-179.

Third, as discussed in [1.1] and [1.8], allocations of bandwidth for dynamic channels are in “units of columns,” each corresponding to 72 bits of information. Ex.1005, 12:6-7. Ex.1003, ¶180.

Accordingly, before making a request for a decrease in the dynamic channel allocation, a node would first compare the bits of data that have been released to the number of bits in a column (“*a predetermined threshold*”) to determine whether a decrease of an entire column (or entire columns) is warranted. Doing so ensures that the decrease in the dynamic channel bandwidth allocation is attributed only to the released sub-channel, and thus will not affect other sub-channels that may be using shared columns of the dynamic frame. Ex.1003, ¶181.

Fourth, as discussed in [1.8], in response to a client request (REQ) to deallocate bandwidth, the arbiter responds by setting a grant section (GNT) in the dynamic frame specifying the total amount of bandwidth granted to the dynamic channel. Ex.1005, 10:35-38. A POSITA would have understood that when the request for decreased bandwidth is granted by the arbiter, the bandwidth (e.g., the

number of columns) allocated to the client node is decreased by the requested amount. Ex.1003, ¶182.

A POSITA would have further understood that the requested amount may also be “*the predetermined quantum*” by which the allocation was previously increased as discussed in [1.8]. When the bursty nature of data traffic has subsided, the increased allocation of resources would have no longer been needed. In this scenario, the additional resources allocated for sending the bursty data on the dynamic channel may be deallocated. Ex.1003, ¶¶183-184.

Therefore, Kovvali renders this limitation obvious. Ex.1003, ¶¶174, 185.

15. Claim 15

[15.0] *A communication network, comprising:*

See analysis at [1.0].

[15.1] *a plurality of nodes, comprising at least first and second nodes; a plurality of links, interconnecting the nodes to provide communications therebetween:*

See analysis at [1.0].

[15.2] *and a dispatcher, which is adapted to allocate a respective number of quanta of one or more resources in the network to each of the nodes,*

See analysis at [1.1] and [11.0].

[15.3] wherein the first node is adapted, upon receiving a request to use a portion of the resources so as to carry the data flow between the first and second nodes over one of a plurality of paths therebetween, each such path comprising a respective sequence of the links,

See analysis at [1.2].

[15.4] to direct the data flow from the first node to the second node over one of the paths without requesting additional resources if the resources already allocated to the first node on the one of the paths are sufficient,

See analysis at [1.3].

[15.5] and to request an increase in an allocation of the requested resources to the first node if the resources already allocated to the first node are insufficient,

See analysis at [1.4] and [1.5].

[15.6] and wherein in response to the increase requested by the first node, the dispatcher is adapted to determine, for the links comprised in each of the paths, respective levels of use of the requested resources due to communications in progress over the network,

See analysis at [1.6].

[15.7] and to select which of the paths is to carry the data flow responsive to the determined levels of use of the requested resources on the links comprised in each of the paths,

See analysis at [1.7].

[15.8] and to increase the allocation of the resources to the first node on the selected path by a predetermined quantum.

See analysis at [1.8].

16. Claim 16

[16.0] *A network according to claim 15, wherein the dispatcher is adapted to select for carrying the data flow the one of the paths that has the lowest level of a predetermined measure of use of the requested resources:*

See analysis at [10.0].

17. Claim 17

[17.0] *A network according to claim 16, wherein the measure of use comprises, for each of at least two of the paths, a maximum level of use of at least one of the requested resources, taken over all of the links comprised in the paths.*

As discussed in [10.0], Kalman provides examples where cost is based on the level of congestion (or another measure of resource use) alone. Referring to Kalman's Figure 3 (discussed in [1.6]), for a given packet/flow between two nodes (e.g. Node 0 to Node 5), in an exemplary scenario there may be congestion on the counterclockwise path (blue) between Node 0 and Node 7, corresponding to a link cost of 10, and no congestion on the remaining links on the counterclockwise path (i.e. Node 7 to Node 6, and Node 6 to Node 5). On the clockwise path (orange), there may be congestion between Node 1 and Node 2, corresponding to a link cost of 25, and no congestion on the remaining links of the clockwise path (i.e. Node 0 to Node 1, Node 1 to Node 2, Node 2 to Node 3, Node 3 to Node 4, and Node 4 to Node 5). Ex.1003, ¶¶197-198.




FIG. 3

Ex.1006, Fig. 3 (annotated); Ex.1003, ¶198.

Applying Kalman's teachings, the counterclockwise path would be selected because the maximum link cost on that path (e.g. 10) is lower than the maximum link cost on the clockwise path (e.g. 25), resulting in the total path cost for the counterclockwise path (blue) being the lower than the path cost for the clockwise path. Ex.1003, ¶199.

Therefore, Kalman renders this limitation obvious. Ex.1003, ¶¶196, 200.

18. Claim 18


[18.0] *A network according to claim 15, wherein the first node is adapted, upon release of the requested resources, to compare an amount of the resources that have been released to a predetermined threshold, and if the amount is greater than the predetermined threshold, to deallocate the predetermined quantum from the first node.*

See analysis at [14.0].

D. Ground 2

1. Masters

Like Kalman, Masters relates to choosing one of a plurality of routes between nodes based on total cost. Ex.1007, 3:34-37. Figure 7 illustrates an example of a network where messages may be routed across various communication paths between nodes, labeled C-G. Ex.1007, 13:57-67; Ex.1003, ¶202.



Ex.1007, Fig. 7 (annotated); Ex.1003, ¶202.

In Figure 7, Masters depicts two possible routes for sending a message to site E from site C, “which reflect taking different routes having different intermediate hops.” Ex.1007, 14:32-37. One possible route (blue) “travels from node C to F to E,” and another route (orange) travels “from nodes C to F to G to E.” Ex.1007, 14:37-41; Ex.1003, ¶203.

“Each potential route is composed of one or more connectors specified in a sequential order. The message is passed sequentially from one connector to the next connector in the route until the message reaches its destination site. Each connector provides direct connectivity between two sites and has an associated

individual cost.” Ex.1007, 3:40-45. In Figure 7, the first exemplary route (blue) from node C to node E uses connectors Vcf and Vfe, and the second exemplary route (orange) uses connectors Vcf, Vfg, and Vge. Ex.1007, 14:33-34, Fig. 8; Ex.1003, ¶204.

Masters describes that a “repository” (including a routing table of message routing information) is used to identify the total cost associated with each of the respective plurality of routes between connectors. Ex.1007, 3:27-39, 4:63-5:9. The route chosen is based on its associated cost, for example, “the lowest cost route.” Ex.1007, 3:27-39, 7:29-38. “The total cost of a route is equal to the individual costs of its constituent connectors, plus an optional increment amount for each intermediate site encountered along the route.” Ex.1007, 3:45-47. Figure 8 provides an example of a routing table with the total cost of various routes from node C to nodes D, E, F, and G. For example, the total cost associated with sending a message along the first exemplary route (blue) using connectors Vcf and Vfe is 13. The total cost associated with sending a message along the second exemplary route (orange) from node C to E using connectors Vcf, Vfg, and Vge is 16. Ex.1003, ¶205.

Site C's Routing Table

<u>Destination</u>	<u>Route</u>	<u>Total Cost</u>
D	Vcd	1
D	Vcf, Vfe, Ved	21
D	Vce, Ved	108
E	Vcd, Vde	12
E	Vcf, Vfe	13
E	Vcf, Vfg, Vge	16
F	Vcf	9
F	Vcd, Vde, Vef	18
F	Vce, Vef	106
G	Vcf, Vfg	12
G	Vcf, Vfe, Veg	18
G	Vcd, Vde, Vef, Vfg	21
G	Vce, Veg	105
Internet	Vcf, Vfg, Vgi	14
.	.	.
.	.	.
.	.	.

Ex.1007, Fig. 8 (annotated); Ex.1003, ¶205.

The value assigned to cost may be “any other measure of resources” including the amount of “processing time” that will be expended if the site connector is used. Ex.1007, 6:2-8. As another example, the cost may reflect the “processing power used to send a message by that connector.” Ex.1007, 16:37-38. “One of the plurality of routes is chosen based on its associated total cost, for example, the lowest cost route, and the message is sent from the originating site to the destination site using the chosen route.” Ex.1007, 3:34-37; Ex.1003, ¶206.

2. Reasons to Combine

A POSITA would have been motivated to combine Masters with Kovvali and Kalman, because Masters also relates to efficient transmission of data in a communication network. Ex.1007, 15:7-10 (“to determine the best route for a message to take to reach a final destination” based on “cost information”); *see also* Section X.C.3 (describing reasons to combine Kovvali and Kalman). Ex.1003, ¶207.

A POSITA would have been motivated to combine Masters with Kovvali and Kalman, and doing so would be within a POSITA’s skillset. The combination would have been obvious because it represents merely the use of Masters’ known technique for selecting a preferred route based on the cost of “processing power” to improve Kalman’s total cost calculation in the same way. When processing power of a node affects the total cost of selecting a path for transmitting data, a POSITA would look to Masters for details on how to account for the cost of processing power. Ex.1003, ¶208.

The combination would have also been obvious because the application of Masters’ technique (selecting a preferred route on a network based on total cost that accounts for “processing power”) to improve Kalman’s method (selecting a preferred path on a ring network based on total cost) ready for improvement and would yield predictable results. Here, the improvement provided to Kalman’s total

cost calculation is accounting the cost of “processing power” that may impact the total cost of a particular path. Ex.1003, ¶209.

It was known in the art that “the processing power available at a switch or router is limited.” Ex.1010, 3:16-17. While that processing power is typically “dedicated to the primary business of the switch/router—switching or routing,” other obligations can sometimes cause “performance drops.” Ex.1010, 3:18-20. Furthermore, the same level of processing power may not be available for all ports on a switch or router. Ex.1010, 3:22-24 (“Putting significant processing power on all the ports of a switch or router is not cost-effective.”). As a result, it would have been obvious to a POSITA for Kovvali’s arbiter node to take into account the available processing power of the nodes in the network when evaluating potential paths for network traffic flowing between two nodes. Ex.1003, ¶210.

A POSITA would have had a reasonable expectation of success in combining the teachings of Masters and Kalman. Kalman describes that the computation of cost associated with sending data between nodes “may take into account a variety of factors.” Ex.1006, 7:66-8:1. Masters simply contemplates that the cost of a route may include the “processing power” used by the connectors to send the data. Ex.1007, 16:33-38. Masters also explicitly suggests the combination by stating that its techniques “find applicability in any network that transmits information between different nodes, or locations.” Ex.1007, 3:65-67. A POSITA

would have understood that Masters’ technique for including “processing power” in a total cost calculation for selecting a preferred path for sending data between nodes would have been readily applicable to the combined network of Kovvali and Kalman. Ex.1003, ¶211.

3. Claim 6

[6.0] *A method according to claim 1, wherein the one or more resources comprise a processing power associated with each of the links.*

First, as discussed in [1.6], Kovvali teaches a BLSR network where each pair of nodes are connected with two unidirectional optical links traveling in opposite directions. Ex.1005, 2:10-16. Further, as discussed in [1.6], Kalman discloses a technique for selecting a preferred direction for transmitting data between nodes based on the direction with the lowest cost. Ex.1006, 7:50-63; Ex.1003, ¶213.

Second, Masters also describes a technique for choosing one of “a plurality of routes for reaching the destination site...based on its associated total cost, for example, the lowest cost route.” Ex.1007, 3:32-36. “Each potential route is composed of one or more connectors specified in a sequential order. The message is passed sequentially from one connector to the next connector in the route until the message reaches its destination site. Each connector provides direct connectivity between two sites and has an associated individual cost.” Ex.1007,

3:40-45; Ex.1003, ¶¶214-215.

Masters describes that the “value assigned to the cost parameter ... may correspond to any other measure of resources--for example, the amount of processing time or bandwidth that will be expended if the site connector is used.” Ex.1007, 6:2-8. In another example, Masters describes that “the individual cost associated with a particular connector represents a ... measure of processing power used to send a message by that connector.” Ex.1007, 16:30-31, 16:37-38; Ex.1003, ¶216.

Accordingly, it would have been obvious to a POSITA that when the cost of a link between two connectors includes “the processing power used to send a message by that connector,” the processing power is “*associated with*” the link directly connecting two sites. Ex.1003, ¶217.

Masters further describes a “repository” that includes “a routing table which lists potential routes for reaching remote sites.” Ex.1007, 4:66-5:5. “The repository includes an identity of at least one destination site, a plurality of routes for reaching the destination site, and a total cost associated with each of the respective plurality of routes. One of the plurality of routes is chosen based on its associated total cost, for example, the lowest cost route....” Ex.1007, 3:31-36; Ex.1003, ¶218.

A POSITA would have found it obvious that calculating the total cost for each of a plurality of routes based on “the processing power used to send a

message” by the respective connectors in those routes, and selecting the “lowest cost route,” renders obvious determining respective levels of resource usage (as discussed above in [1.6]) and selecting a path responsive to the determined levels of use (as discussed in [1.7]) based on “*a processing power associated with each of the links.*” Ex.1003, ¶¶212, 219.

E. Ground 3

1. RFC2892

RFC2892 specifies a “Spatial Reuse Protocol” (SRP) for use with ring based media. Ex.1018, 1. RFC2892 describes a “SRP Fairness Algorithm (SRP-fa)” for allocating bandwidth fairly around the ring. Ex.1018, 9. When there is no congestion, “[e]ach node can send traffic up to max_usage. Max_usage is a per node parameter that limits the maximum amount of low priority traffic a node can send.” Ex.1018, 29; Ex.1003, ¶220.

2. Reasons to Combine

A POSITA would have been motivated to combine RFC2892 with Kovvali and Kalman, because RFC2892 also relates to “efficient use of bandwidth” in a ring network. Ex.1018, 1. As discussed in [4.0], Kovvali and Kalman each describe “spatial reuse” for transmitting traffic around the ring. In particular, Kovvali describes explicitly a “Spatial Reuse Protocol (SRP)” allowing reuse of ring bandwidth across different sections of a ring of nodes. Ex.1005, 4:17-18. Kovvali

describes further that “[a] supplementary algorithm called SRP Fairness algorithm (SRP-fa) was also defined to achieve fair sharing of bandwidth between various nodes in a ring topology.” Ex.1005, 4:18-21. A POSITA carrying out the teachings of SRP and SRP-fa in Kovvali would look to known documents such as RFC2892 that specify the SRP and SRP-fa respectively. Ex.1003, ¶221.

3. Claim 7

[7.0] *A method according to claim 1, wherein selecting which of the paths is to carry the data flow comprises comparing an amount of the one or more resources requested to a resource budget assigned to the first node, and permitting the data flow only if allocating the requested resources will not cause a total of the resources allocated to the first node to exceed the budget.*

First, as discussed in [1.1], the arbiter node in Kovvali assigns “dynamic channels” to particular nodes and adjusts their allocated rates in response to “requests” from client nodes. Ex.1005, 9:1-11. “At any time, each dynamic channel is allocated zero or more columns” of the elastic section in the dynamic frame. Ex.1005, 10:3-5; Ex.1003, ¶223.

Second, RFC2892 describes a “SRP Fairness Algorithm (SRP-fa)” that “governs access to the ring” and that sets parameters that “limit the maximum amount of low priority traffic a node can send.” Ex.1018, 9, 28-29; Ex.1003, ¶224. In particular, RFC2892 describes a variable “max_usage” defined as “a per node parameter that limits the maximum amount of low priority traffic a node can send.”

Ex.1018, 29. “Each node can send up to max_usage.” Ex.1018, 29; Ex.1003, ¶224.

In another example, RFC2892 describes “SRP-fa pseudo-code” that defines a similar variable “MAX_ALLOWANCE” as the “configurable value for max allowed usage for this node.” Ex.1018, 30; Ex.1003, ¶225. At each clock cycle, the amount of data transmitted by a host is compared to “MAX_ALLOWANCE” to check for the condition that the amount of data transmitted by the host (i.e. “my_usage”) is less than “MAX_ALLOWANCE,” or specifically “my_usage < MAX_ALLOWANCE.” Ex.1018, 31; Ex.1003, ¶226. Only if this condition is true is it “OK to send host packets.” Ex.1018, 31; Ex.1003, ¶226.

Accordingly, a POSITA would have thus found it obvious that “max_usage” corresponds to a “*resource budget assigned to the first node*” because a node can only send “up to” max_usage which is the value for the “maximum amount of low priority traffic” a node can send. Similarly, a POSITA would have found it obvious that “MAX_ALLOWANCE” corresponds to a “*resource budget assigned to the first node*” because it is the value configured for the “max[imum] allowed usage for this node.” Ex.1003, ¶227.

Third, as discussed in [1.5], Kovvali describes that a client node may “request additional bandwidth for the corresponding bursty channel from the arbiter.” Ex.1003, ¶228. It would have been obvious to a POSITA that an arbiter would compare the total bandwidth requested for a dynamic channel with the

maximum allowed usage for the node (e.g. “max_usage” or “MAX_ALLOWANCE”) so that the arbiter does not grant bandwidth that exceeds a maximum amount (e.g., the maximum number of dynamic channel columns) allowable for the node. It would have been obvious to a POSITA for Kovvali’s arbiter not to allocate more bandwidth than the maximum value allowed so that network bandwidth can be allocated fairly, as discussed in RFC2892 relating to the SRP Fairness Algorithm. Ex.1003, ¶229.

Therefore Kovvali renders this limitation obvious. Ex.1003, ¶¶222, 230.

4. Claim 8

[8.0] *A method according to claim 7, wherein comparing the amount of the one or more resources comprises comparing the amount of each of the resources requested to the resource budget assigned for each of the resources, and wherein permitting the data flow comprises permitting the flow only if all of the resources requested for at least one of the paths are within the budget.*

As discussed in [7.0], Kovvali and RFC2892 teach comparing the total bandwidth requested for a dynamic channel (“*each of the resources requested*”) with the maximum allowed bandwidth usage (e.g. “max_usage” or “MAX_ALLOWANCE”) for the node (“*resource budget assigned for each of the resources*”). Ex.1003, ¶232. A POSITA would have found it obvious that when the resources requested for transmitting a dynamic channel exceeds the maximum allowed bandwidth usage, the arbiter would either deny the request (e.g. GNT

00000) or return a lower value up to “[a]ll available columns” (e.g. GNT 11111) for the channel in the frame so that the maximum number of columns that the node’s bandwidth usage does not exceed the configured maximum value (e.g. “max_usage” or “MAX_ALLOWANCE”). Ex.1003, ¶232.

A POSITA would have also found it obvious that if a dynamic channel had a minimum data requirement that exceeds the maximum allowed bandwidth usage for the node (“*the budget*”), the arbiter would deny the request for resources. Ex.1003, ¶233. Denying the request would ensure the network only delivers an expected quality of service according to the minimum data requirement of the dynamic channel. Ex.1003, ¶233.

It would have been obvious to a POSITA for Kovvali’s arbiter node to take into account quality of service requirements of different data flows because it was well known in the art that network traffic may have “different QoS [quality of service]” parameters with varying bandwidth requirements. Ex.1020, 2:42-49 (“initializing a QoS manager at an admit limit (AL) representing a bandwidth utilization capacity allocated to a service offered on the weakest link of the network; receiving a service request including a service type and a required bandwidth; comparing the AL with the required bandwidth and accepting the request whenever the required bandwidth does not exceed the AL, and rejecting the request whenever the required bandwidth exceeds the AL”). Accordingly, Kovvali

and RFC2892 render obvious “*permitting the flow only if all of the resources requested for at least one of the paths are within the budget.*” Ex.1003, ¶234.

Therefore, Kovvali, Kalman, and RFC2892 render this limitation obvious. Ex.1003, ¶¶231, 235.

XI. DISCRETIONARY DENIAL IS INAPPROPRIATE

A. No basis for 35 U.S.C. § 325(d) denial.

The combinations of references presented in the grounds below have not previously been presented to the Office. Neither Kovvali nor Masters were cited or considered by the Examiner during prosecution of the '279 patent. *See* Ex.1002. While RFC2892 and Kalman were before the Examiner, they were never cited in any rejection, nor even mentioned in any Office action. *See* Ex.1002, 88, 195. Petitioner’s analysis above shows that the Examiner erred by not rejecting the claims based in part on Kalman and RFC 2892, and therefore discretionary denial is inappropriate under the second prong of *Becton, Dickinson*. *See* IPR2017-01586, Paper 8 at 17-18. Additionally, denial under 35 U.S.C. § 325(d) is not warranted because all the challenges presented in this petition are based in part on Kovvali, which (1) teaches pre-allocating bandwidth so that transmission may occur “*without requesting additional resources*”—the feature argued to gain allowance in prosecution, Ex.1002, 190—and (2) is neither cumulative nor redundant to the prior art considered in prosecution. *Advanced Bionics, LLC v. MED-EL*

Elektromedizinische Geräte GmbH, IPR2019-01469, Paper 6 (PTAB Feb. 13, 2020) (precedential).

B. Discretionary denial under *Fintiv* is not appropriate.

The Board balances six “*Fintiv*” factors in considering denial under § 314. *Apple Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 (Mar. 20, 2020) (precedential). These factors strongly favor institution. The district court case is in its very early stages. Patent Owner filed its complaint on September 26, 2023, Ex.1011, and served it on September 27, 2023. Patent Owner provided its disclosure of asserted claims and infringement contentions on December 22, 2023. Ex.1012. Petitioner has since worked diligently to prepare this Petition, which is being filed well within the one-year timeframe allowed by Congress, and before invalidity contentions are due on February 19, 2024, Ex.1113.

1. Whether evidence exists that a stay may be granted if a proceeding is instituted

A motion to stay has not yet been filed and the Board should not infer or attempt to predict the outcome if such a motion is later filed. *Sand Revolution II LLC v. Continental Intermodal Group – Trucking LLC*, IPR2019-01393, Paper 24 at 7 (June 16, 2020) (informative); *see also Dish Network L.L.C. v. Broadband iTV, Inc.*, IPR2020-01359, Paper 15 (Feb. 12, 2021) (“It would be improper to speculate, at this stage, what the Texas court might do regarding a motion to

stay...”). Thus, this factor is neutral.

2. Proximity of the court’s trial date to the board’s projected statutory deadline for a final written decision

The Docket Control Order sets December 4, 2024, as the date for the claim construction hearing, and June 2, 2025, as the trial date. Ex.1014, 1, 4. The expected date for a Final Written Decision in this case is July 2025.

This factor weighs against discretionary denial because the Petitioner has worked expeditiously to prepare and file this petition approximately one month after receiving Patent Owner’s infringement contentions on December 22, 2023. Ex.1010, 4; *Fintiv*, Paper 11 at 11–12.

Furthermore, while trial is currently scheduled to begin in June 2025, “scheduled trial dates are often subject to change.” *In re Apple Inc.*, 979 F.3d 1332, 1344, n.5 (Fed. Cir. 2020). In *Sand*, the “relative close proximity” of a trial date scheduled **5 months** before the FWD deadline “weigh[ed] marginally in favor of *not* exercising discretion to deny institution” given the district court’s own recognition of a “continuing degree of ... uncertainty of the court’s schedule.” *Sand* at 8-10. Here, even if the EDTX trial commences against Petitioner in June 2025, the FWD in this case is expected just one month later in July 2025, which is well within the 5 month time period in *Sand*.

Thus, this factor weighs against discretionary denial.

3. Investment in the parallel proceeding

The co-pending litigation is in its early stages, and the investment in it has been minimal. Fact discovery is not set to close until January 13, 2025, and expert discovery is not set to close until February 18, 2025. Ex.1014, 3. The Docket Control Order sets December 4, 2024, as the date for the *Markman* hearing. Ex.1014, 4. Thus, this factor weighs against discretionary denial. *PEAG LLC v. Varta Microbattery GMBH*, IPR2020-01214, Paper 8 at 17 (Jan. 6, 2021) (little investment in the parallel proceeding weighed against discretionary denial).

4. Overlap between issues raised in the petition and in the parallel proceeding

There will be no overlap of prior art issues. If the Board institutes trial, Petitioner will cease asserting in the co-pending litigation the combination of references on which trial is instituted for the claims, to the extent Petitioner even asserts the same combination in the district court. This factor weighs against discretionary denial. *See Verizon v. Huawei*, IPR2020-01079, Paper 10 at 38 (Jan. 14, 2021) (finding a similar stipulation “mitigates the concern about overlapping issues” and weighs “against discretionary denial of the Petition.”).

5. Whether the petitioner is a defendant

Petitioner is a defendant in the litigation. Ex.1011. This is true of most Petitioners in IPR proceedings. While the *Fintiv* case indicates that a difference

between the district court defendant and the petitioner may weigh against discretionary denial, nothing within the *Fintiv* case suggests that the same party between the proceedings weighs in favor of discretionary denial. *See HP Inc. v. Slingshot Printing LLC*, IPR2020-01084, Paper 13 at 9 (having the “same parties as parallel proceeding” makes factor 5 “neutral”). Accordingly, this factor is neutral and should not be a basis for denying institution.

6. Other circumstances that impact the Board’s exercise of discretion, including the merits

As discussed in detail above, the prior art presented in this Petition renders the Challenged Claims unpatentable as obvious. In view of the striking similarities between the ’279 patent claims and Kovvali, Kalman, Masters, and RFC2892, this petition “presents compelling evidence of unpatentability,” and thus this factor weighs against discretionary denial. Ex.1021, 2 (“PTAB will not rely on the *Fintiv* factors to discretionarily deny institution in view of parallel district court litigation where a petition presents compelling evidence of unpatentability.”). In particular, the ’279 patent claims were argued in prosecution as being allowable because they provide for data transmissions “*without requesting additional resources.*” Ex.1002, 189-190. Kovvali teaches an “initial data rate allocation” that, when sufficient, allows transmissions to occur “*without requesting additional resources.*” Ex.1005, 12:5-7 (“initial data rate allocations... in units of columns of the dynamic frame”);

see supra, [1.1], [1.3]. Requests for additional bandwidth are made “from time to time,” such as when a “bursty channel” requires more bandwidth. Ex.1005, 10:35-38, 11:10-13 (“from time to time... the bandwidth allocated to an existing channel is requested to be changed”). If the data being transmitted is accommodated by the existing bandwidth allocation, then the data is transmitted using that allocation and without requesting additional resources. Ex.1005, 12:15-27. Kovvali’s teaching of the concepts argued to gain allowance in prosecution—together with the cited prior art teaching all remaining claim limitations—shows the compelling merits of this petition.

As such, because the *Fintiv* factors are either neutral or weigh against discretionary denial. Because this Petition was filed approximately eight months before the statutory bar date, institution should not be denied on discretionary factors.

C. Discretionary denial under *General Plastic* is not appropriate

The ’279 patent has not been challenged in any prior IPR petition, so none of *General Plastic* discretionary institution factors apply to this Petition. *See General Plastic Indus. Co., Ltd. v. Canon Kabushiki Kaisha*, IPR2016-01357, Paper 19 at 16 (Sept. 6, 2016) (Section II.B.4.i. precedential).

XII. CONCLUSION

Petitioner has established a reasonable likelihood that the Challenged Claims

are unpatentable.

Respectfully submitted,

Dated: January 26, 2024
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, Texas 75201
Customer No. 27683

/Theodore M. Foster/
Theodore M. Foster
Lead Counsel for Petitioner
Registration No. 57,456

XIII. MANDATORY NOTICES

A. Real party-in-interest

Pursuant to 37 C.F.R. § 42.8(b)(1), Petitioner certifies that the real party-in-interest is Cisco Systems, Inc.

B. Related matters

Pursuant to 37 C.F.R. § 42.8(b)(2), to the best knowledge of the Petitioner, the '279 patent is or was involved in the following cases:

Case Heading	Number	Court	Filed
<i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i>	2:23-cv-0041	EDTX	Sept. 26, 2023

C. Lead and back-up counsel and service information

Lead Counsel

Theodore M. Foster
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, TX 75201

Phone: (303) 382-6205
Fax: (214) 200-0853
ipr.theo.foster@haynesboone.com
USPTO Reg. No. 57,456

Back-up Counsel

Brian C. Kwok
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, TX 75201

Phone: (415) 293-8916
Fax: (214) 200-0853
brian.kwok.ipr@haynesboone.com
USPTO Reg. No. 58,828

David L. McCombs
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, TX 75201

Phone: (972) 739-8636
Fax: (214) 200-0853
david.mccombs.ipr@haynesboone.com
USPTO Reg. No. 32,271

Please address all correspondence in this proceeding to lead and back-up counsel. Petitioner consents to service in this proceeding by email at the addresses above, on condition that Patent Owner similarly consents to service in this proceeding by email.

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioner hereby certifies, in accordance with and reliance on the word count provided by the word-processing system used to prepare this Petition, that the number of words in this paper is 13,966. Pursuant to 37 C.F.R. § 42.24(d), this word count excludes the table of contents, table of authorities, mandatory notices under § 42.8, certificate of service, certificate of word count, appendix of exhibits, and any claim listing.

Dated: January 26, 2024

/Theodore M. Foster/
Theodore M. Foster
Lead Counsel for Petitioner
Registration No. 57,456

CERTIFICATE OF SERVICE

The undersigned certifies that, in accordance with 37 C.F.R. § 42.6(e) and 37 C.F.R. § 42.105, service was made on Patent Owner as detailed below.

Date of service January 26, 2024

Manner of service Priority Mail Express International

Documents served Petition for *Inter Partes* Review Under 35 U.S.C. § 312 and 37 C.F.R. § 42.104 of U.S. 7,035,279;
Petitioner's Power of Attorney;
Petitioner's Exhibit List;
Exhibits 1001-1021.

Persons served May Patents Ltd.
c/o Dorit Shem-Tov
P.O.B 7230
Ramat-Gan, - 5217102
ISRAEL

/Theodore M. Foster /
Theodore M. Foster
Lead Counsel for Petitioner
Registration No. 57,456