

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner,

v.

PORTSMOUTH NETWORK CORP.,
Patent Owner.

IPR2024-00503
Patent No. 7,035,279

PETITIONER REPLY

TABLE OF CONTENTS

Table of Contents	2
Petitioner’s Exhibit List	3
I. Introduction.....	5
II. Status of Copending Litigation.....	5
III. Claim Construction.....	5
A. Terms from district court proceeding.....	5
1. “ <i>quantum</i> ”	6
2. “ <i>predetermined</i> ”	8
B. “ <i>resources</i> ”	11
IV. Portsmouth does not address many of the reasons to combine Kovvali and Kalman.....	12
V. Portsmouth argues the references individually under the guise of alleging hindsight bias.....	13
VI. Alleged “fundamental differences” would not have prevented the combination of Kovvali and Kalman.	16
VII. Kovvali renders obvious “receiving a request at the first node...”	17
VIII. Kovvali and Kalman render the Challenged Claims obvious under Petitioner’s proposed district court claim constructions	24
IX. Kovvali renders obvious Claims 7-8	27
A. Portsmouth does not address many of the reasons to combine Kovvali and Kalman with RFC2892.....	27
B. The combination of Kovvali and Kalman with RFC2892 render obvious a “ <i>resource budget</i> .”	30
X. Conclusion	31
Certificate of Service	33

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	Curriculum Vitae 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, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Sept. 26, 2023)
Ex.1012	Notice of Disclosures, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Jan. 8, 2024)
Ex.1013	Order, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Dec. 19, 2023)
Ex.1014	Docket Control Order, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , 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. et al., 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
Ex.1022 (NEW)	Tsiang, D. et al., INTERNET DRAFT, “The SRP MAC Layer Protocol” (Jun. 10, 1999), <i>available at</i> https://datatracker.ietf.org/doc/html/draft-tsiang-srp-00
Ex.1023 (NEW)	Plaintiff’s Opening Claim Construction Brief, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441, Dkt. 78 (E.D. Tex. Oct. 23, 2024)
Ex.1024 (NEW)	Declaration of Dr. James Olivier in Support of Portsmouth’s Opening Claim Construction Positions, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441, Dkt. 78-1 (E.D. Tex. Oct. 23, 2024)
Ex.1025 (NEW)	U.S. 6,272,131 (“Ofek”)
Ex.1026 (NEW)	Transcript of Remote Videotaped Deposition of James Olivier, Ph.D. (Jan. 16, 2025)
Ex.1027 (NEW)	U.S. 5,461,611 (“Drake”)
Ex.1028 (NEW)	Defendant Cisco’s Responsive Claim Construction Brief, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441, Dkt. 83 (E.D. Tex. Nov. 6, 2024)
Ex.1029 (NEW)	Declaration of James Olivier, Ph.D., <i>Metacluster LT, UAB v. Bright Data Ltd.</i> , IPR2022-00687, Ex.1002.
Ex.1030 (NEW)	Declaration of James Olivier, Ph.D., <i>Roku, Inc. v. Flexiworld Techs., Inc.</i> , IPR2022-01114, Ex.1002.

I. Introduction

The Petition shows that claims 1-18 (the “Challenged Claims”) of the ’279 patent are unpatentable. Patent Owner’s Response (“POR”) fails to distinguish the claims over the prior art, leaves unrebutted several motivations to combine references, and presents meritless strawman defenses that mispresent the arguments presented in the Petition.

The POR therefore fails to overcome the Petition’s evidence and arguments. The Board should find the Challenged Claims unpatentable.

II. Status of Copending Litigation

The complaint filed against Petitioner Cisco Systems, Inc. was dismissed without prejudice on December 27, 2024. The *Markman* hearing that had been scheduled for January 7, 2025 was cancelled, and there is no longer any active copending litigation involving the ’279 patent.

III. Claim Construction

A. Terms from district court proceeding

Patent Owner Portsmouth Network Corp. notes that the claim terms “*predetermined*” and “*quantum*” were proposed for construction in the copending district court action. POR, 4. Portsmouth further alleges that “Petitioner has failed to address these constructions....” POR, 4. But Cisco Systems, Inc. filed the Petition in this case on January 26, 2024, while the parties did not even exchange

proposed claim terms for construction in the district court action until six months later. Ex.1014, 5 (“July 31, 2024... Exchange Proposed Claim Terms”). Thus, this Petitioner Reply is Cisco’s first opportunity to address in this proceeding the claim construction issues that have developed over the past year.

For the purposes of this proceeding, Cisco agrees that the Board does not need to formally construe the terms proposed for the district court case. *See* POR, 4. Portsmouth does not argue that Cisco’s proposed constructions would distinguish the claims over the prior art. To the contrary, Portsmouth’s position in district court was that Cisco’s proposed constructions were “too narrow.” Ex.1023, 7; *see also* Ex.1024, ¶¶30, 33 (“Cisco’s proposed construction... is unnecessarily narrow....”).

To the extent that the Board determines that construing “*predetermined*” or “*quantum*” would be helpful, the Board should adopt the constructions proposed by Cisco in the district court proceeding:

1. “*quantum*”

If construction is necessary, Cisco proposes construing “*quantum*” as “an amount of a requested resource by which to increase or decrease.” This interpretation is based on the specification and credits a key aspect of the claimed invention—that “[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.

A patent tribunal will deviate from the plain and ordinary meaning of a term “when a patentee sets out a definition and acts as his own lexicographer.” *Golden Bridge Tech. v. Apple Inc.*, 758 F.3d 1362, 1365 (Fed. Cir. 2014). Here, the ’279 patent defines “node_quantum” as an “Amount by which to increase or decrease node_used_In or Out at each resource allocation or deallocation.” Ex.1001, 5:26–28. Thus, the specification effectively defines “quantum” as “an amount of a requested resource by which to increase or decrease.”

The claims themselves further demonstrate that quantum is “an amount of a requested resource by which to increase or decrease,” and not just any amount. The first step of the “method for routing a data flow,” in Claim 1, is “allocating a respective number of quanta of one or more resources in the network to each of the nodes.”¹ The specification explains that “[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. In other words, when a node is allocated resources in quanta (i.e., increments of quantum as opposed to an exact amount needed), the node is allocated an extra portion of unused resources (the difference between the allocated quanta and the amount needed). The node can utilize the unused resources for another data flow without requesting additional resources.

¹ The parties agree that “quanta” is the plural of “quantum.” Ex.2007, 2.

Ex.1001, cl. 1 (“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”). As a result, the resource requests occur less frequently, lowering the burden on the system. *See* Ex.1001, cl. 1 (“requesting an increase in an allocation to the first node” only in the event “the resources already allocated to the first node are insufficient”), 5:38-39.

Applicants emphasized that “allocation of resources in quanta of this sort is useful in limiting the frequency” of allocation operations and “reduce[s] the communication and computing burden.” Ex.1002, 119-120, 190. Applicants pointedly distinguished prior art that allocated “the exact bandwidth required.” Ex.1002, 190. Applicants explained that the claims allow resources left over from previously allocated quanta to be used without requesting additional resources. Ex.1002, 190. If the claimed technology allowed resources to be allocated in any arbitrary amount, then allocations would be made exactly—and the distinctions and benefits argued in prosecution would have been illusory.

Thus, a *quantum* is “an amount of a requested resource by which to increase or decrease.” Ex.1028, 9-13.

2. “*predetermined*”


If construction is necessary, Cisco’s proposes construing “*predetermined*” as

“determined before the data flow request.” This construction is supported by the ’279 patent claims and specification, and it clarifies *before what* a quantum must be determined. Ex.1028, 13-17.

Claim 1 recites that the first step of its “method for routing a data flow” is “*allocating a respective number of quanta of one or more resources in the network to each of the nodes,*” followed by “*receiving a request at the first node to . . . carry the data flow*” (i.e., the data flow request), and then, if the available resources are insufficient, “requesting an increase in an allocation to the first node of the requested resources . . . [and] *in response to [that] request, . . . increasing the allocation* of resources to the first node on the selected path by a predetermined quantum.” Ex.1001, cl. 1. This claim language demonstrates two facts. First, the node had been allocated quanta (at step 1) **before** any data flow request was made (at step 2), confirming that the size and type of an allocable quantum is known **before** the data flow request. Second, the last step of Claim 1 uses the term “predetermined quantum” to refer back to the same type and size of quanta that were allocated in the first step of Claim 1 (quanta that are determined before the data flow request).

The ’279 patent’s specification confirms that initial setup of the network occurs before the data flow request and includes configuring the quantum, further supporting Cisco’s construction. Figure 3 (shown in color below) depicts the initial

setup of the claimed network. **First**, “[i]n preparation for this method, a dispatcher entity is chosen or assigned . . . at an assignment step 30” (green). Ex.1001, 5:59–63. **Second**, “[a]t a learning step 32, [the] dispatcher [] learns the topology . . . and builds the tables of network information to be used subsequently in flow allocations” (yellow). Ex.1001, 6:6–10. Notably, “*the dispatcher is configured with* the following parameters for each of the nodes with respect to each of the resources . . . *node_quantum*—Amount by which to increase or decrease node_used_In or Out at each resource allocation or deallocation.” Ex.1001, 5:14–28. Indeed, even “each node [] maintains its own record of the . . . node_quantum.” Ex.1001, 5:33-50. Thus, the quantum is predetermined and known by both the dispatcher and requesting node at the step where the dispatcher learns the topology (in yellow). **Third**, after that predetermination, the “node receives a request from one of its clients to send data to another node” in step 34 (red), i.e., a data flow request is made. Ex.1001, 6:11–14.



Ex.1001, Fig. 3 (annotated)

Thus, Figure 3 and its surrounding description—like the claims—confirm that the quantum is determined *before* the data flow request.

For these reasons, to the extent Board believes a construction for this term is necessary, it should adopt Cisco’s construction for “*predetermined*.” Ex.1028, 13-17.

B. “resources”

Portsmouth raises arguments relating to the claimed step of “*receiving a request... to use a portion of the resources*” that implicate the meaning of the word “*resources*.” See, e.g., POR, 17 (“receiving, processing data and then multiplexing it to another node are operations of data transfer, not requesting resources”). The ’279 patent provides an express definition of the claim term “*resources*”:

The term “resources” in this context, and throughout the

specification and claims of the present patent application, refers not only to bandwidth, but also to processing power used to transmit, receive, encrypt or decrypt data, as well as to 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.

To the extent that the Board finds that the plain and ordinary meaning of “*resources*” is inadequate, the Board should apply this definition, at a minimum, in evaluating Portsmouth’s arguments that are directly contrary to the ’279 patent’s express definition.

IV. Portsmouth does not address many of the reasons to combine Kovvali and Kalman.

The Petition described in detail many discrete reasons that a POSITA would have had to combine the ideas described in Kovvali and Kalman. *See, e.g.*, Paper 2 at 23-26, 46-47. Rather than address these reasons, the POR identifies differences between Kovvali’s and Kalman’s networks. While nominally framed as being about “hindsight” bias, the POR’s arguments are just improper attacks against the references individually. *See, e.g.*, POR, 12 (“Kalman teaches a distributed routing algorithm.... Kovvali, on the other hand, teaches a central arbiter...”). Left un rebutted are the Petition’s proffered reasons to combine:

1. 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.”
Petition, 23.

2. The combination would have been obvious because the application of Kalman’s technique to improve Kovvali’s method “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.”
Petition, 23-24.
3. 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.” Petition, 25-26.
4. A POSITA would have a reasonable expectation of success “because of the overall similarities between the network architectures and technologies described in Kovvali and Kalman” (e.g. bi-directional ring networks transmitting frames between nodes over an optical network), and Kalman describes that its path-selection technique may be applied “to any network transmitting frames, cells, or using any other protocol.” Petition, 26.

Any one of these reasons is sufficient to establish the obviousness of combining Kovvali and Kalman. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 420 (2007) (“any need or problem known in the field and addressed by the patent can provide a reason for combining the elements in the manner claimed.”).

V. Portsmouth argues the references individually under the guise of alleging hindsight bias.

Portsmouth argues that the Petition’s arguments relating to improved efficiency improperly rely on “hindsight.” Specifically, Portsmouth argues that

“increased efficiency” was “the problem already articulated in the ’279 patent.”

POR, 12. The Board should not be misled by these arguments for many reasons.

First, Portsmouth’s argument implies that the ’279 patent’s named inventors were the first to recognize the value of improving efficiency—a position so strained that it borders on being comical. Improving efficiency is such a well-known goal that courts formally recognize it as a perennial motivation for combining prior art references. *See Intel Corp. v. Qualcomm Inc.*, 21, F.4th 784, 797 (Fed. Cir. 2021) (“[B]ecause such [efficiency] improvements are technology-independent, universal, and even common-sensical, there exists in these situations a motivation to combine prior art references even absent any hint of suggestion in the references themselves.” (emphasis in original, cleaned up, citing *KSR Int’l Co. v. Teleflex Inc.*, 550, U.S. 398, 421 (2007))).

Second, Portsmouth provides no citation to where the ’279 patent supposedly recognized that “increased efficiency” was needed or otherwise any kind of “problem” – apparently because there is no support. The only mention of efficiency (or any related word) in the specification is in the field of invention, which neither mentions any problems with the prior art nor suggests that efficiency was a distinguishing characteristic of the claimed subject matter. Ex.1001, 1:6-9.

The ’279 patent’s background section suggests a variety of issues that it seeks to address—none of them referencing efficiency. The patent identifies issues

like “how IP flows should be assigned to high priority levels,” “how load control should be accomplished,” and “how the [routing] costs are to be determined” or “updat[ed]... automatically, while the network is running, in response to network resource constraints.” Ex.1001, 1:36-67. Portsmouth’s “hindsight” argument is illusory: the ’279 patent could not have been used as a roadmap to improve efficiency in the prior art because it simply does not provide any meaningful discussion of efficiency.

Third, Portsmouth’s “hindsight” argument is just another instance of (improperly) attempting to distinguish the prior art references from each other, not from the claims: “Kovvali’s and Kalman’s methods of managing communications over ring networks are quite different.” POR, 13. As the Board is well-aware, “one cannot show non-obviousness by attacking references individually.” *In re Keller*, 642 F.2d 413, 426 (C.C.P.A. 1981).

Finally, Petition’s motivation-to-combine analysis was firmly grounded in the prior art. The Petition articulated reasons why a POSITA considering Kovvali would have found it obvious—without any knowledge of the ’279 patent—to make various modifications, such as modifying Kovvali’s arbiter to choose which direction to send a particular flow. *See* Petition, 46-47. Portsmouth does not dispute that in a bi-directional-ring implementation of Kovvali, there would be “two path options to choose from,” e.g., clockwise and counterclockwise. Pet., 42. Kalman

describes—in the same bi-directional-ring context—known techniques for making that decision. Pet., 43-45. Their combination is merely the use of a known technique to improve a similar system in the same way. Pet, 23-26; *Intel Corp. v. PACT XPP Schweiz AG*, 61 F.4th 1373, 1380-1381 (Fed. Cir. 2023).

VI. Alleged “fundamental differences” would not have prevented the combination of Kovvali and Kalman.

Portsmouth argues that “a POSITA would not combine Kovvali with Kalman because of their fundamental differences.” POR, 13. Specifically, Portsmouth argues that “Kalman’s structure is not compatible with the network structure disclosed in Kovvali.” POR, 13-14.

Like the previous hindsight bias argument, this is a strawman argument that ignores the combination proposed in the Petition. The Petition did not argue for superimposing the entirety of “Kalman’s structure” onto Kovvali’s network. Rather, the proposed combination is for Kovvali’s arbiter to adopt some of the congestion-control and path-selection techniques described in Kalman, such as choosing between clockwise and counterclockwise paths based on the level of traffic congestion in each direction. *See* Petition at 42-47. The only structural prerequisite is for Kovvali’s network to support propagating messages in either direction around the ring—which Kovvali’s “**bi-directional** line switched ring” does. Ex.1005, 8:33-36; Ex.1003, ¶¶84-85; Petition at 35-36.

Portsmouth argues that Kalman’s path-selection techniques could not be

adapted to work in Kovvali's network because "the two references are different and are directed to different outcomes, i.e., efficient switchovers during a node failure vs. accommodating unpredictable data flows." POR, 14. This is conclusory attorney argument, not even supported by testimony from Portsmouth's declarant. Portsmouth fails to develop the argument or explain why or how path selection differs when responding to an unexpected node failure vs. to an unpredictable data flow. Notably, both Kovvali and Kalman describe making dynamic path selections, allowing them to respond to changing conditions regardless of cause. Ex.1005, 9:65-10:1, 8:47-53; Ex.1006, 16:10-13.

VII. Kovvali renders obvious "receiving a request at the first node..."

Regarding the Petition's claim language analysis, Portsmouth argues that Kovvali's arbiter node corresponds to the claimed "*first node*" because the arbiter node receives a bandwidth allocation request. POR, 15-16. But that is not how the Petition mapped the claim language to Kovvali's disclosure. Instead, the Petition identified "Node B or Node C" as the "*first node*." *See, e.g.*, Petition, 35. Thus, Portsmouth's argument about the arbiter "*receiving a request*" attacks a strawman position and is without merit.

In attempting to respond to the Petition's actual mapping, Portsmouth argues that "receiving data at a node" is not "a request" for the node to use its resources to forward the data. POR, 17. This argument fails to appreciate the nature and

significance of the data received by Kovvali's Node B or Node C from local dynamic channels. As Kovvali explains, these local dynamic channels are used to carry "bursty" data that may have "time-varying data rate requirements." Ex.1005, 4:30-34, 8:47-53. The data packets received by Node B or Node C present dynamic and unpredictable demands on network resources, both at the nodes themselves and on the links between nodes. Ex. 1003, ¶83.

Kovvali teaches that Node B or Node C perform various functions to process and multiplex the data received from local dynamic channels onto the ring network. Ex.1005, 14:27-37. These functions include "queuing, scheduling and policing processes" that are necessary to handle the outgoing packet traffic. Ex.1005, 14:29-33. But performance of these steps is not guaranteed, as such packets are transmitted as "best effort data traffic." Ex.1005, Abstract. It was well-known that "best effort" packets may be discarded, that is, not transmitted to their destinations. Ex.1025, 1:60-66 (in "best effort" service "variations [in] traffic intensity often lead to congestion that results in excessive delays and loss of packets"); Petition, 58 (citing Ex.1009, 5:65; Ex.1008, 3:53-57). Kalman expressly notes the possibility that "lower priority" traffic may be "dropped." Ex.1006, 3:31-34. Thus, every "best effort" packet received at Kovvali's Node B or Node C seeks to use some network resources whose availability is not guaranteed. Transmitting a best-effort packet will necessarily consume some resources, such as transmission

processing power and communication bandwidth. *See* Ex.1001, 2:12-18

(“resources... refers not only to bandwidth, but also to **processing power used to transmit, receive, encrypt or decrypt data, as well as to any other allowable factor that may affect the ability of a network link to convey a data flow** between nodes.”).

Consistent with Cisco’s argument, Portsmouth’s expert, Dr. Olivier, testified that the term “resources” as used in the claims is “kind of broad” and “could be other things [such as] link bandwidth, link utilization” and buffer “queues.” Ex.1026, 12:5-17. Dr. Olivier also agreed that if a first node is transmitting packets to a second node, the first node uses both “network bandwidth” and its own “processing power” to forward packets. Ex.1026, 20:19-21:16. And in the context of “best effort” transmission he admitted that if a best effort packet was received by a node, “you’d queue up the best effort traffic until you got time to send it across a link ...” Ex.1026, 24:11-17. Forwarding Kovvali’s bursty traffic unambiguously requires the use of “*resources*,” which may or may not be available. *See* Ex.1003, ¶¶71-72. Each data packet of bursty traffic seeks to use those resources.

The Petition’s mapping of an individual packet as an implicit “*request...to use... resources*” is consistent with how Applicants argued the claims in prosecution. *See* Petition, 35. Applicants repeatedly distinguished the “Drake”

prior art (Ex.1027) because it required a “reservation of facilities for each new data stream.” Ex.1002, 119, 190. Avoiding explicit requests was Applicants’ alleged innovation. Consistent with that argument, Applicants amended the claims to change originally presented claim language of “receiving a **request to allocate** one or more resources...” to the ultimately allowed language of “*receiving a request... to use a portion of the resources*”:

1. (Currently amended) 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:

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

receiving a request at the first node to allocate use a portion of the one or more resources in the network so as to carry the data flow between the first and second ~~ones~~ of the nodes over one of a plurality of paths therebetween, each such path comprising a respective sequence of the links;

Ex.1002, 112 (highlighting added).

Applicants emphasized that claim 1 is directed to a node that can use its own previously allocated resources to carry a new data flow (i.e. “without requesting additional resources”):

By contrast, in Drake's system, the nodes must request a reservation of facilities for each new data stream. Drake states explicitly that "the quality of service management of a local area network is almost entirely concentrated in a single QoS management station... QoS allocations are maintained only so long as a requestor continues to refresh the request" (col. 2, line 65 – col. 3, line 10). Drake's nodes do not have a prior allocation of quanta that they can use to direct data flows in the manner recited in claim 1. When Drake's QoS management station (identified in Fig. 1 as allocator station 20) allocates a QoS transmission path for a multimedia file, it determines the exact bandwidth required and subtracts the bandwidth from each component on the path (col. 11, lines 11-21). Drake does not teach or suggest the possibility that a node could use its own allocated quanta of resources to carry a data flow, without requesting additional resources, if the allocated resources are sufficient.

Ex.1002, 190 (emphasis added, discussing Drake (Ex.1027)).

The specification and dependent claims also support Petitioner's mapping of a packet-to-be-transmitted to an implicit "*request.. to use... resources so as to carry the data flow*" by stating that the claimed *data flow* is conveyed immediately, as part of the *selecting* step:

Most preferably, selecting the first or the second path includes conveying the data flow over an inner or outer data link ring within the network provided by a Spatial Reuse Protocol (SRP).

Ex.1001, 2:62-3:4; *see also* claim 4.

The *selecting* step – which includes conveying the data flow – is, notably, performed "*in response to the request*" that was received (e.g., a packet in the data flow). Ex.1001, 9:33. There is no further step recited of "receiving" the data flow

after the “*request... to use... resources.*”

A POSITA would have understood that the receipt of data from local dynamic channels by Node B or Node C for transmission to another node is indeed a “*receiving a request... to use a portion of the resources*” on the ring network. Receiving a packet causes the node to use (if available) the resources required to transmit the packet toward its destination. Ex. 1003, ¶85. This is consistent with the claim language of [1.2], which does not require an explicit or formal request, but rather is open to any form of a request. Thus, Kovvali discloses or renders obvious this limitation. Ex. 1003, ¶86.

Portsmouth argues that treating a packet as a “*request*” to use network resources would mean that “Node B or Node C would be receiving millions of requests per second.” POR, 17. But Portsmouth fails to explain how this is an impediment to obviousness. The claims do not limit how often *resources* may be requested or used.

Portsmouth also argues that Kovvali’s Node B or Node C cannot “receiv[e] the request... to use a portion of the resources... to carry the data flow” because Kovvali’s arbiter would have already assigned some resources to those nodes. POR, 18. This argument conflicts with other language in claim 1, whose first step expressly recites “*allocating... resources... to each of the nodes.*” Plainly, the prior allocation of resources is not an impediment to a node later “*receiving a request...*

to use a portion of the resources.” Similarly, claim 1 further recites responding to the “*request*” by testing “*if the resources already allocated... are sufficient.*” Again, the mere fact that some resources are already allocated to handle packets does not mean that a packet seeking to use those resources cannot be a “*request... to use a portion of the resources.*”

Portsmouth also argues that “handling traffic is completely different from **requesting** traffic.” POR, 18. That argument is inscrutable, and in any event, fails to address the claim language, which does not recite “requesting traffic.”

Finally, Portsmouth argues that Kovvali is distinguished because Nodes B and C are “assigned to carry the data flow by the arbiter component in Kovvali.” POR, 18. Portsmouth argues that since Kovvali’s Nodes are “assigned” bandwidth for carrying their respective dynamic channels, the Nodes cannot receive a “*request*” to use that assigned bandwidth. This argument appears to confuse the “*receiving a request*” step with the earlier step of “*allocating... resources... to each of the nodes.*” See Petition, 27-32. Again, Portsmouth ignores that Applicants’ alleged innovation was a system that pre-allocates resources, and they amended the claim language from a “request to **allocate**” to a “*request... to use.*” See Ex.1002, 119, 112.

In summary, Portsmouth fails to distinguish Kovvali’s bandwidth management scheme, which operates like the claimed system:

- nodes are pre-allocated with network bandwidth to be used at their discretion (Ex.1001, 4:46-55; Ex.1005, 12:1-6);
- when a node receives data from a client to transmit data to another node, if the pre-allocated bandwidth is sufficient, the node uses that pre-allocated bandwidth to transmit the data (Ex.1001, 6:48-51, 6:54-57; Ex.1005, 6:4-14); and
- if the already allocated bandwidth is insufficient, the node submits a request to the arbiter node for additional resources (Ex.1001, 6:57-60; Ex.1005, 10:33-38).

VIII. Kovvali and Kalman render the Challenged Claims obvious under Petitioner’s proposed district court claim constructions

The terms “*predetermined*” and “*quantum*” appear in the phrase “*a predetermined quantum*” in independent claims [1.8] and [15.8], which are each rendered obvious by Kovvali:

[1.8] *and increasing the allocation of the resources to the first node on the selected path by a predetermined quantum*

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

Kovvali teaches that a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:33-38; Ex.1003, ¶121. In particular, the client node sets a request section (REQ) within the dynamic frame overhead to specify the amount of requested bandwidth.

Ex.1005, 10:33-38; Ex.1003, ¶121.

Kovvali describes that the request (REQ) can be for a certain number of columns (amount) by which to allocate (increase) or deallocate (decrease) to the dynamic channel. Ex.1005, 10:33-38, Ex.1003, ¶121. As shown in the table below, Kovvali provides four request REQ values to allocate a specified number of additional columns (i.e. 100, 101, 110, and 111) and three REQ values to deallocate a specified number of columns (i.e. 001, 010, and 011) of bandwidth:

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

The arbiter in Kovvali responds to the request by setting a grant section (GNT) within the dynamic frame specifying the amount of bandwidth granted. Ex.1005, 10:33-38, Ex.1003, ¶121. 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” (a binary value corresponding to a total allocation of 6 columns). The resources allocated to the client node for sending data on the dynamic channel would have been increased by 1 column of additional bandwidth. Ex.1003, ¶121. Accordingly, it would have been obvious to a POSITA that when the arbiter grants the requested bandwidth change, the number of columns (amount) allocated to the client node would be increased (or decreased) by the requested amount (“*quantum*”). Pet. 49; Ex.1003, ¶123.

Kovvali also teaches that the number of columns that may be requested (“*quantum*”) by a node is encoded in a table. Ex.1005, 10:39-61. In particular, Kovvali discloses that certain values of request REQ values correspond to known numbers of columns to allocate (or deallocate) to a dynamic channel. Ex.1005, 10:46-61. For example, a REQ value of 100 means to allocate “1 additional column” and a REQ value of 111 means to allocate “10 additional columns.” Ex.1005, 10:52, 55, Ex.1003, ¶123. Since Kovvali’s disclosure specifies the sizes of the available allocation adjustments, it would have been obvious that each of the request REQ values in Kovvali’s table are “determined before the data flow request.” Ex.1003, ¶122; *see also id.* ¶123 (“the allocations are for predetermined, quantized amounts of bandwidth.”). A client node requests an allocation change using the number of columns encoded in the table, “thereby reducing the number

of bits needed for the control channel.” Ex.1005, 10:39-40; Ex.1003, ¶122.

Therefore, Kovvali’s teaching of requesting certain numbers of columns already encoded in a table to be allocated or deallocated to a dynamic channel renders obvious “*a predetermined quantum.*” Pet. 48-50, Ex.1003, ¶124.

IX. Kovvali renders obvious Claims 7-8

A. Portsmouth does not address many of the reasons to combine Kovvali and Kalman with RFC2892

Portsmouth argues inaccurately that Petitioner provides “no tangible reason for a POSITA to combine” Kovvali and Kalman with RFC2892. POR, 19. That threadbare response simply ignores most of Dr. Mir’s direct testimony on this issue. *Compare* POR, 19 & Ex.2004, ¶53 to Pet., 75-76 & Ex.1003, ¶221 (explaining how “Kovvali and Kalman each describe ‘spatial reuse’” techniques standardized in RFC2892).

The Petition provides many discrete reasons that a POSITA would have combined RFC2892 with Kovvali and Kalman that Portsmouth leaves un rebutted (Petition, 75-76):

1. Kovvali and Kalman each describe “spatial reuse” for ring traffic. In particular, Kovvali describes explicitly using a “Spatial Reuse Protocol (SRP)” allowing reuse of ring bandwidth across different sections of a ring of nodes. Ex.1005, 4:12-18.
2. Kovvali also describes explicitly using a supplementary Spatial Reuse Protocol algorithm “called SRP Fairness algorithm (SRP-fa) ... defined to achieve fair sharing of bandwidth between various nodes in a ring topology.” Ex.1005, 4:18-21.

3. RFC2892 is prior art publication that specifies the “Spatial Reuse Protocol” (SRP) for use with ring based media. Ex.1018, 1.

During his deposition, Dr. Olivier agreed that both Kovvali and RFC2892 both refer to “Spatial Reuse Protocol (SRP)” and “SRP fairness algorithm (SRP-fa).” Ex.1026, 55:15-22. And Dr. Olivier testified previously that “RFCs (and similar standards documents) posted on the Internet are published in the ordinary course by established standards organizations, and are **intended to be viewed by the interested Internet engineering audience.**” Ex.1029, ¶81 (emphasis added). Based on that testimony, RFC2892 was intended to be considered by the individuals interested in Kovvali’s spatial reuse concepts. Dr. Olivier has also affirmed that those of skill in the art were motivated to refer to RFC documents to “ensure interoperability” among networking devices. Ex.1030, ¶65.

In deposition, Dr. Olivier pivoted to identifying a new alleged source of “confusion” that he did not articulate in his Declaration: whether Kovvali refers to the same version of SRP and SRP-fa that were promulgated in RFC2892. Ex.1026, 55:3-6, 15-17. According to Dr. Olivier, “it would seem more logical that Kovvali was referring to Version 1” of SRP and SRP-fa (Ex.1026, 55:7-10), instead of the “second version of the protocol (V2)” described in RFC2892. Ex.1018, 1. In short, Dr. Olivier was merely pointing out that RFC2892 is dated August 2000, which is after Kovvali was filed in March 2000. Ex.1026, 55:15-17. Since both documents are prior art to the ’279 patent, their relative timing is simply irrelevant. *See* 35

U.S.C. § 103(a) (“...the subject matter as a whole would have been obvious **at the time the invention...**”); *Siemens Med. Solutions U.S. Inc. v. Saint-Gobain Ceramics & Plastics Inc.*, 637 F.3d 1269,1293 (Fed. Cir. 2011) (Prost, J., dissenting) (“Obviousness is assessed at the time of the invention...”).

Dr. Olivier does not deny that a version of the Spatial Reuse Protocol existed before Kovvali was filed. During his deposition, Dr. Olivier was provided an Internet-Draft version of an Internet Engineering Task Force (IETF) document, dated June 10, 1999 (thus, predating Kovvali), bearing the title “The SRP MAC Layer Protocol.” Ex.1026, 71:13-15; *see* Ex.1022. That earlier draft also includes descriptions of the “Spatial Reuse Protocol (SRP)” and “SRP Fairness Algorithm title (“The SRP MAC Layer Protocol”) and lists the exact same authors (“D. Tsiang” and “G. Suwala”). Ex.1026, 72:18-73:10.

Finally, Dr. Olivier also does not argue that Kovvali’s reference to SRP or SRP-fa is referring to some other type of spatial reuse protocol, other than that specified by RFC2892 and its predecessor drafts. In fact, Dr. Olivier testified that he was “not... aware of any other spatial reuse protocol standards.” Ex.1026, 51:2-4.

In summary, Dr. Olivier’s professed “confusion” is a self-induced, meaningless distraction. **As of the relevant priority date** (January 9, 2001), a POSITA considering Kovvali and Kalman would have looked to known documents

for further information about SRP and SRP-fa, including the then-current published standard in RFC2892 (Ex.1018). Ex.1001, 75-76; Ex.1003, ¶¶221, 229.

B. The combination of Kovvali and Kalman with RFC2892 render obvious a “*resource budget*.”

Portsmouth further argues that the “parameters described by the RFC2892 reference are not aligned with the node budget described in claims 7 and 8.” POR, 19. Specifically, Portsmouth argues that “max_usage” is a “per node parameter” that “*cannot* be used to limit data flows which are not low priority traffic.” POR, 20-21. Portsmouth fails to explain this reasoning, and its declarant merely repeats these statements verbatim. Ex.2004, ¶55.

Portsmouth’s argument that “max_usage” is a “per node parameter” is misdirected, since the claim recites “*a resource budget assigned to the first node.*” See Petition, 77.

Portsmouth’s attempt to distinguish the prior art based on traffic priority levels is similarly misdirected. Portsmouth’s argument fails to address the prior art’s teaching that “max_usage” **is used** to “limit[] the maximum amount of low priority traffic a node can send.” Ex.1018, 29. Regarding “MAX_ALLOWANCE,” Portsmouth acknowledges that the parameter is “used in controlling low priority traffic.” POR, 21. Since claim 7 does not recite that the “*resource budget*” pertains to high-priority traffic, it is simply irrelevant whether “max_usage” and “MAX_ALLOWANCE” pertain to only certain traffic priorities.

Finally, Portsmouth argues that max_usage and MAX_ALLOWANCE are “congestion management and internode communication” parameters, whereas the ’279 patent describes “resource management parameters.” POR, 21. But this argument is contradicted by Dr. Olivier’s own testimony where he admitted that the ’279 patent is at least “inferentially” managing network congestion because “congestion is something that the network provider considers something they need to plan for perhaps in the future.” Ex.1026, 35:21-36:14. Moreover, Portsmouth fails to tether this supposed distinction between parameters to the claim language, which does not recite “parameters” at all. Instead, claims 7 and 8 refer to a “*resource budget*.” Portsmouth does not distinguish the prior art’s example parameters used to manage bandwidth (a “*resource*,” Ex.1001, 2:12-18, claim 5) from the claimed “*resource budget*.”

X. Conclusion

The Board should find the Challenged Claims unpatentable.

Respectfully submitted,

Dated: February 4, 2025
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

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 Petitioner's Reply, that the number of words in this paper is 5,555. This word count excludes the table of contents, certificate of service, certificate of word count, signature block, and appendix of exhibits. *See* 37 C.F.R. §42.24(c).

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

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. § 42.6(e), the undersigned certifies that a true and correct copy of the foregoing was served on the Patent Owner as detailed below:

<i>Date of service</i>	February 4, 2025
<i>Manner of service</i>	Electronic Email: jnuttall@steptoe.com
<i>Documents served</i>	Petitioner Reply; Petitioner's Exhibit List; and Exhibits Ex.1022 – Ex.1030
<i>Persons Served</i>	James R. Nuttall STEPTOE LLP 227 West Monroe, Suite 4700 Chicago, Illinois 60606 Katherine D. Cappaert Heather P. Hildreth STEPTOE LLP 1330 Connecticut Avenue N.W. Washington, D.C. 20036

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