IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

ORCKIT CORPORATION, Plaintiff, v. CISCO SYSTEMS, INC., Defendant.	Civil Action No. 2:22-cv-276 **JURY TRIAL DEMANDED**

DEFENDANT'S MOTION FOR A STAY PENDING
***INTER PARTES* REVIEW PROCEEDINGS ON ALL FOUR ASSERTED PATENTS**

TABLE OF CONTENTS

Page

I. INTRODUCTION 1

II. FACTUAL BACKGROUND 2

A. The Present Litigation Is At An Early Stage 2

B. Cisco's Pending IPRs Challenge All Asserted Claims On Grounds Distinct From Those Considered By The Patent Examiners 3

III. LEGAL STANDARD 4

IV. ARGUMENT 6

A. A Stay Will Not Prejudice Orckit. 6

B. The Early Stage Of The Case Weighs In Favor Of A Stay. 8

C. A Stay Will Simplify This Case 9

D. At Minimum, The Court Should Deny This Motion Without Prejudice So That Cisco Can Renew The Motion After The IPR Institution Decisions. 11

V. CONCLUSION 12

TABLE OF AUTHORITIES

Page(s)

CASES

Alcohol Monitoring Sys., Inc. v. ActSoft, Inc.,
Nos. 07-cv-02261-PAB, 08-cv-01226, 2011 WL 5075619 (D. Colo. Oct. 25, 2011)..... 7

Anascape, Ltd. V. Microsoft Corp.,
475 F. Supp. 2d 612 (E.D. Tex. 2007)..... 10

Arbor Glob. Strategies LLC v. Samsung Elecs. Co.,
No. 2:19-cv-00333-JRG-RSP, 2021 WL 66531 (E.D. Tex. Jan. 7, 2021) 11

Aylus Networks, Inc. v. Apple Inc.,
856 F.3d 1353 (Fed. Cir. 2017)..... 10

Clinton v. Jones,
520 U.S. 681 (1997)..... 4

Customedia Techs., LLC v. DISH Network Corp.,
No. 2:16-cv-129-JRG, 2017 WL 3836123 (E.D. Tex. Aug. 9, 2017) 4

Cywee Grp. Ltd. v. Samsung Elecs. Co.,
No. 2:17-cv-00140-WCB-RSP, 2019 WL 11023976 (E.D. Tex. February 14, 2019) 5, 7, 8

Ethicon LLC v. Intuitive Surgical, Inc.,
2019 WL 1276029 (D. Del. Mar. 20, 2019) 10

Ethicon, Inc. v. Quigg,
849 F.2d 1422 (Fed. Cir. 1988)..... 4

e-Watch Inc. v. Apple, Inc.,
No. 2:13-cv-1061-JRG-RSP, 2015 WL 12915668 (E.D. Tex. Mar. 25, 2015) 11

Finjan, Inc. v. Symantec Corp.,
139 F. Supp. 3d 1032 (N.D. Cal. 2015) 5

In re Intel Corp.,
No. 2021-168, 2021 WL 4427875 (Fed. Cir. 2021) 11

Kove IO, Inc. v. Amazon Web Servs., Inc.,
No. 18-cv-8175, 2022 WL 683666 (N.D. Ill. Mar. 8, 2022) 7

Meetrix IP, LLC v. Zoho Corp.,
No. 1:22-cv-588-LY (E.D. Tex. February 28, 2023)..... 5, 7

Microlinc, LLC v. Intel Corp.,
No. 2:07-cv-488 TJW, 2010 WL 3766655 (E.D. Tex. Sept. 20, 2010)..... 6

Murata Mach. USA v. Daifuku Co.,
830 F.3d 1357 (Fed. Cir. 2016) 5

NFC Tech. LLC v. HTC Am., Inc.,
No. 2:13-cv-1058-WCB, 2015 WL 1069111 (E.D. Tex. Mar. 11, 2015) 5, 10

Norman IP Holdings, LLC v. TP-Link Techs., Co.,
No. 6:13-cv-384, 2014 WL 5035718 (E.D. Tex. Oct. 8, 2014) 8

Onpoint Sys., LLC v. Protect Animals With Satellites, LLC,
No. 4:20-cv-657, 2022 WL 2704166 (E.D. Tex. July 12, 2022) 4, 5, 9

Ramot at Tel Aviv Univ. Ltd. v. Cisco Sys., Inc.,
No. 2:19-cv-00225-JRG, 2021 WL 121154 (E.D. Tex. Jan. 13, 2021) 11

Regents of the Univ. of Minnesota v. LSI Corp.,
926 F.3d 1327 (Fed. Cir. 2019) 4

Spine Holdings, LLC v. Orthofix Medical, Inc.,
No. 4:20-cv-77-SDJ (E.D. Tex. June 8, 2020) 5

Uniloc 2017 LLC v. LG Elecs. U.S.A., Inc.,
No. 3:18-cv-3071-N, 2020 WL 374545 (N.D. Tex. Jan 23, 2020) 6

Uniloc USA, Inc. v. Samsung Elecs. Am., Inc.,
No. 2:16-cv-642, 2017 WL 9885168 (E.D. Tex. June 13, 2017) 9

Village Green Techs. LLC, v. Samsung Elecs. Co. Ltd.,
No. 2:22-cv-00099-JRG, 2023 WL 416419 (E.D. Tex. Jan. 25, 2023) 6, 8

VirtualAgility Inc. v. Salesforce.com, Inc.,
759 F.3d 1307 (Fed. Cir. 2014) 9

Wi-LAN, Inc. v. LG Elecs., Inc.,
No. 3:17-cv-00358, 2018 WL 2392161 (S.D. Cal. May 22, 2018) 5

STATUTES

35 U.S.C. § 315(e)(2) 11

OTHER AUTHORITIES

H.R. Rep. 112-98, pt. I (2011) 4

I. INTRODUCTION

The Court should stay this case because Defendant Cisco Systems, Inc. ("Cisco") filed IPR petitions on every asserted claim in this case. Cisco's filings come early in the case—before any depositions have taken place and more than a month before the parties exchange claim terms—and stand to create the ultimate issue simplification, entirely mooting the case. In such circumstances, each factor weighs in favor of staying the case.

First, Orckit faces no undue prejudice because it is a non-practicing entity that does not compete with Cisco.

Second, the case is at a very early stage. The claim construction process has not begun, the Markman hearing is six months away, fact discovery just started and does not close for seven months, opening expert reports are not due for seven months, Plaintiff has only produced thirty-nine documents to date, no disputed issues have been decided by this Court, and trial is nearly a year away. Additionally, Cisco's motion timing is not strategic. The same day that Cisco filed its last IPR, it informed Orckit that Cisco planned on seeking a stay and requested a meet and confer at Orckit's first availability. Cisco filed this motion the day after the parties met and conferred.

Third, because Cisco's IPRs cover all asserted claims, resolution of Cisco's IPRs will significantly narrow, or entirely moot, this litigation. A stay would avoid the risk of proceeding with a likely unnecessary and burdensome litigation—including claim construction, fact discovery, expert discovery, summary judgment, pre-trial filings, and trial. And even if some of the challenged claims were to survive the IPRs, a stay will simplify validity issues in this case. Such simplifications are particularly important given the incredibly complex nature of this case; Orckit accused Cisco of infringing over 100 claims across four unrelated patents implicating hundreds of accused products and identified more than a dozen foreign witnesses in its disclosures.

To the extent the Court believes that granting a stay would be premature before the Patent Office issues its IPR institution decisions, Cisco respectfully requests that the Court withhold ruling on Cisco's motion until such institution decision or deny Cisco's motion without prejudice, to refile after the Patent Office issues its institution decisions.

II. FACTUAL BACKGROUND

A. The Present Litigation Is At An Early Stage

On July 22, 2022 Orckit filed its first complaint accusing Cisco of infringing the four Patents-in-suit, which together included over 130 claims.[1] Orckit only asserted one "exemplary" claim per Patent-in-suit and did not identify any allegedly infringing features of any of the products it accused of infringing. Dkt. No. 1. Cisco responded by filing a Motion to Dismiss Plaintiff's Complaint For Patent Infringement Pursuant to Fed. R. Civ. P. 12(b)(6). Dkt. No. 15. Orckit then filed an amended complaint on October 14, 2022, which for the first time identified some accused features of a few product lines, but still only identified one claim per Patent-in-suit. Dkt. No. 21.

On November 3, 2022 Orckit served infringement contentions identifying for the first time the 104 claims it accuses Cisco of infringing. Cisco promptly objected to the adequacy of those infringement contentions and explained to Orckit that Orckit's infringement contentions did not adequately provide notice of its infringement theories as required by this Court's rules. After Cisco identified the deficiencies in Orckit's contentions, the parties began meeting and conferring about amending those contentions without the need for judicial intervention through November and December of 2023. Orckit served finalized amended contentions, providing the first actual notice of its infringement theories, on January 19, 2023.

[1] Orckit has asserted U.S. Patent Nos. 6,680,904; 7,545,740; 8,830,821; and 10652,111 (collectively, the "Patents-in-suit"). Cisco filed petitions for *Inter Partes* review on all asserted claims of all of the Patents-in-suit (collectively, the "Co-Pending IPRs").

Cisco filed its first two IPR petitions on January 9, 2023, only 6 weeks after Cisco first answered Orckit's complaint (Dkt. No. 26) and before Orckit finalized its infringement contentions.

On February 2, 2023, Cisco served its initial invalidity contentions.

Cisco filed its third IPR petition six weeks after the first two, on February 21, 2023, and filed its last petition just three weeks later on March 14, 2023. In February and March 2023, the parties met and conferred regarding alleged deficiencies in Cisco's invalidity contentions. To resolve the issue without judicial intervention, Cisco plans on filing an unopposed motion to amend its invalidity contentions at the end of this month.

Fact discovery in the present matter is in its infancy. Since fact discovery opened, the parties have served, but not responded to, their first sets of interrogatories. Orckit has only produced 39 total documents so far.[2] Neither party has taken any depositions yet. The parties will not exchange proposed claim terms for claim construction until May 4, 2023 and will not have a claim construction hearing until September 7, 2023. The deadline to complete fact discovery is seven months away on October 19, 2023. Expert discovery has not started. Dispositive motions are eight months away, set for November 27, 2023, and the trial is set for just under a year away on March 4, 2024.

B. Cisco's Pending IPRs Challenge All Asserted Claims On Grounds Distinct From Those Considered By The Patent Examiners

By March 14, 2023—prior to the parties finalizing contentions, responding to any interrogatories, taking depositions, and well before the Claim Construction hearing—Cisco filed four IPRs, covering each Asserted Patent and all 104 asserted claims in the present litigation:

[2] Pursuant to P.R. 3-3 and 3-4, Cisco produced tens of thousands of technical documents regarding the accused products and prior art with its invalidity contentions, but Cisco's burden is not relevant to the stay factors as the movant.

IPR No.	Asserted Patent	Filing Date	Claims Challenged	Exhibit
IPR2023-00401	7,545,740	January 9, 2023	1-31 (all claims)	Ex. 1
IPR2023-00402	8,830,821	January 9, 2023	1-20 (all claims)	Ex. 2
IPR2023-00554	10,652,111	February 21, 2023	1-9, 12-24, 27-31 (all asserted claims)	Ex. 3
IPR2023-00714	6,680,904	March 14, 2023	1-26 (all claims)	Ex. 4

The IPRs Cisco filed rely on prior art that was not considered by the Patent Office during the original prosecution of the patent. Indeed, only one of the twelve prior art references that Cisco asserts in an IPR ground even appeared on the face of a challenged patent. *Compare* Exs. 1-4, *with* Patents-in-suit.

III. LEGAL STANDARD

The Court possesses the inherent power to control its own docket, including the power to stay proceedings. *Customedia Techs., LLC v. DISH Network Corp.*, No. 2:16-cv-129-JRG, 2017 WL 3836123, at *1 (E.D. Tex. Aug. 9, 2017) (citing *Clinton v. Jones*, 520 U.S. 681, 706 (1997)); *Ethicon, Inc. v. Quigg*, 849 F.2d 1422, 1426–27 (Fed. Cir. 1988). A goal of an IPR is "to limit unnecessary and counterproductive litigation costs." *Regents of the Univ. of Minnesota v. LSI Corp.*, 926 F.3d 1327, 1335 (Fed. Cir. 2019) (citing H.R. Rep. 112-98, pt. I, at 40 (2011)). "A stay is particularly justified when the outcome of a PTO proceeding is likely to assist the court in determining patent validity or eliminate the need to try infringement issues." *Onpoint Sys., LLC v. Protect Animals With Satellites, LLC*, No. 4:20-cv-657, 2022 WL 2704166, at *1 (E.D. Tex. July 12, 2022) (internal quotations omitted).

When deciding whether to stay a case pending IPR, district courts will consider "(1) whether the stay will unduly prejudice the nonmoving party, (2) whether the proceedings before the court have reached an advanced stage, including whether discovery is complete and a trial date

has been set, and (3) whether the stay will likely result in simplifying the case before the court." *Onpoint*, 2022 WL 2704166 at *2 (quoting *NFC Tech. LLC v. HTC Am., Inc.*, No. 2:13-cv-1058-WCB, 2015 WL 1069111, at *2 (E.D. Tex. Mar. 11, 2015)). "Based on those factors, courts determine whether the benefits of a stay outweigh the inherent costs of postponing resolution of the litigation." *Id.* The Federal Circuit and courts in this District have also considered "whether a stay will reduce the burden of litigation on the parties and the court." *See, e.g.*, *Murata Mach. USA v. Daifuku Co.*, 830 F.3d 1357, 1362 (Fed. Cir. 2016) (internal quotations omitted); *Cywee Grp. Ltd. V. Samsung Elecs. Co.*, No. 2:17-cv-00140-WCB-RSP, 2019 WL 11023976, at *2 (E.D. Tex. February 14, 2019).

District Courts, including in this District, have granted stays prior to institution decisions, particularly when the case is early enough to lead to considerable conservation of resources and when the IPRs will clarify and streamline the issues for the court. *See, e.g.*, *Meetrix IP, LLC v. Zoho Corp.*, No. 1:22-cv-588-LY, Dkt. No. 43 at 3 (E.D. Tex. February 28, 2023) (granting stay prior to institution because all three considered factors weighted in favor of a stay); *Spine Holdings, LLC v. Orthofix Medical, Inc.*, No. 4:20-cv-77-SDJ, Dkt. No. 8 at 2 (E.D. Tex. June 8, 2020) (staying case pending IPR prior to even filing IPR petition when "neither party will be prejudiced and that the case is early enough in the litigation process that a stay is likely to result in considerable conservation of both judicial and party resources."); *Finjan, Inc. v. Symantec Corp.*, 139 F. Supp. 3d 1032, 1037–38 (N.D. Cal. 2015) (granting stay "pending a decision by the PTO concerning whether to institute IPR" and noting "were the Court to deny the stay until a decision on institution is made, the parties and the Court would expend significant resources on issues that could eventually be mooted by the IPR decision"); *Wi-LAN, Inc. v. LG Elecs., Inc.*, No. 3:17-cv-00358, 2018 WL 2392161, at *2 (S.D. Cal. May 22, 2018) (granting a stay "pending the PTO's decisions

regarding institution of [Defendant's] IPR petitions" and finding that a "stay would further promote the interest of justice and judicial economy").

IV. ARGUMENT

The three stay factors—simplification of issues, stage of the proceedings, and potential undue prejudice to the non-moving party—favor granting a stay. ***First***, Orckit is a non-practicing entity seeking monetary damages for alleged infringement by products that Cisco has sold for years. ***Second***, the case is in its infancy; fact discovery has barely started and the vast majority of the work in the case lies in the future. ***Third***, Cisco's IPRs cover all 104 asserted claims of the four Patents-in-suit and stand to largely simplify, if not moot, this entire case.

A. A Stay Will Not Prejudice Orckit.

A stay pending Cisco's IPRs will not unduly prejudice Orckit. Orckit and Cisco are not competitors. Indeed, Orckit was only formed in April 2022, and in its complaint Orckit did not provide a principal place of business nor allege that it sold any products. Dkt. No. 21. Orckit only stated that Orckit Communications Ltd.—an unrelated entity with a similar name—used to create telecommunications infrastructure systems before being liquidated. *Id.* at 4. Because Orckit does not manufacture or sell any products, Orckit cannot allege that it would be harmed by customer losses or by injury to market share during a stay. *See Microlinc, LLC v. Intel Corp.*, No. 2:07-cv-488-TJW, 2010 WL 3766655, at *2 (E.D. Tex. Sept. 20, 2010). "Such a lack of competition weighs against a finding of undue prejudice." *Village Green Techs. LLC, v. Samsung Elecs. Co. Ltd.*, No. 2:22-cv-00099-JRG, 2023 WL 416419, at *2 (E.D. Tex. Jan. 25, 2023) (citing *Uniloc 2017 LLC v. LG Elecs. U.S.A., Inc.*, No. 3:18-cv-3071-N, 2020 WL 374545, at *1 (N.D. Tex. Jan 23, 2020)).

Mere delay, without more, is insufficient to establish undue prejudice. *See Cywee*, 2019 WL 11023976, at *2 . The fact that a stay in this case would be no more than "mere delay" is emphasized by Orckit seeking damages and not preliminary injunctive relief. Dkt. 21 at 22, 30, 38, 49, 50; Ex. 5 (Plaintiff's Supplemental Contentions: Asserted Claims by Product Category (as amended on 1/19/2023)) at 82. Orckit will be able to collect damages for alleged infringement that occurred during the stay. Moreover, one of the accused patents is expired; damages are not even accruing on that patent. *See Kove IO, Inc. v. Amazon Web Servs., Inc.*, No. 18-cv-175, 2022 WL 683666, at *3 (N.D. Ill. Mar. 8, 2022) (finding where the asserted patents are expired, "any prejudice resulting from a delay will not be undue prejudice"). Additionally, the accused products span decades and many of those products are no longer sold by Cisco. *Compare* Ex. 5 *with* Ex. 6 (Cisco list of End-of-Sale and End-of-Life Products) (https://web.archive.org/web/20230307092403/https://www.cisco.com/c/en/us/products/eos-eol-listing.html, Mar. 7, 2023); *see Meetrix*, No. 1:22-cv-588-LY, Dkt. No. 43 at 3 ("[Plaintiff] did not file suit against [Defendant] for years after the accused products at issue were first launched. Therefore, the court concludes that the first factor weighs in favor of a stay.").

Nor will a stay give the Cisco an unwarranted tactical advantage. To the contrary, a stay would prevent either party from taking inconsistent positions in the Patent Office and in this case because it will allow this Court to benefit from reviewing the full record of the parties' claim construction arguments in the IPRs, and the Patent Office's resolution of them, before addressing claim construction here. Thus a stay would prevent tactical advantages rather than create them. *See, e.g.*, *Alcohol Monitoring Sys., Inc. v. ActSoft, Inc.*, Nos. 07-cv-02261-PAB, No. 08-cv-01226, 2011 WL 5075619, at *6 (D. Colo. Oct. 25, 2011) ("Allowing plaintiff to alter its position [from that asserted to the Patent Office] would give plaintiff the unfair advantage of retaining the . . .

patent while pursuing an infringement claim based on a position inconsistent with the prior successful position.").

B. The Early Stage Of The Case Weighs In Favor Of A Stay.

Cisco has diligently pursued its IPR petitions and then a stay in this case. *See Village Green*, 2023 WL 416419, at *3 ("The Court also considers whether the defendant acted with reasonable dispatch in filing its petitions for *inter partes* review and then, after the petitions were granted, in filing its motion for a stay.") (internal quotations omitted). Cisco filed its first two IPR petitions before even receiving Orckit's first amended infringement contentions, requested a meet and confer with Orckit regarding a stay on the same day as filing the last IPR on March 14, 2023, and promptly filed this motion after meeting and conferring with Orckit.

Most of the work for the parties and the Court in this case remains ahead, which favors granting Cisco's request for a stay. *See, e.g.*, *Norman IP Holdings, LLC v. TP-Link Techs., Co.*, No. 6:13-cv-384, 2014 WL 5035718, at *3 (E.D. Tex. Oct. 8, 2014) ("Courts often find the stage of litigation weighs in favor of a stay if there remains a significant amount of work ahead for the parties and the court, even when the parties and/or the court have already devoted substantial resources to the litigation.") (internal citations omitted). The claim construction process has yet to begin, and the claim construction hearing is just over six months away and will occur just within about a week of the first two institution decision deadlines for the Co-Pending IPRs. The case is also still in the early stages of fact discovery, with the close of fact discovery about seven months away. The parties have not deposed a single witness, expert discovery has not begun, summary judgment is eight months away, and trial is about a year away. At this point, "[t]he most burdensome parts of the case . . . all lie in the future." *Cywee*, 2019 WL 11023976, at *6.

With claim construction and the substantial discovery deadline well into the future, the case is at an ideal stage for a stay. *See, e.g.*, *VirtualAgility Inc. v. Salesforce.com, Inc.*, 759 F.3d 1307, 1317 (Fed. Cir. 2014) (determining that a stay pending administrative review is proper where "there remained eight months of fact discovery, the joint claim construction statements had yet to be filed, and jury selection was a year away"). A stay would be especially useful to conserve resources here, where the IPRs have the potential to obviate the need for the District Court case entirely.

C. A Stay Will Simplify This Case.

Granting a stay would simplify this case by potentially eliminating all issues in this litigation. "[T]he most important factor bearing on whether to grant a stay in this case is the prospect that the *inter partes* review proceeding will result in simplification of the issues before the Court." *Onpoint*, 2022 WL 2704166, at *3. "A stay is particularly justified when the outcome of a PTO proceeding is likely to assist the court in determining patent validity or eliminate the need to try infringement issues." *Id.* at *1 (internal quotations omitted). That purpose would be served here, as the pending IPR petitions address all 104 asserted claims in all four Patents-in-suit. If the Court stays this case and the PTAB later invalidates the asserted claims of the Asserted Patents in the IPR proceedings, the Court and the parties will have saved significant costs, time, and resources that they would otherwise expend litigating this case now. And even if only some of the claims are invalidated, "there is a significant likelihood that the outcome of the IPR proceedings will streamline the scope of this case to an appreciable extent." *Uniloc USA, Inc. v. Samsung Elecs. Am., Inc.*, No. 2:16-cv-642, 2017 WL 9885168, at *1 (E.D. Tex. June 13, 2017); *see Village Green*, 2023 WL 416419, at *6 ("[S]hould the IPRs result in the cancelation of some

or all of the asserted claims, 'either some portion of the litigation will fall away, or the litigation will come to an end altogether.'") (quoting *NFC Tech.*, 2015 WL 1069111, at *4).

Moreover, regardless of the outcome, statements made during the IPR proceedings will very likely narrow issues of infringement and invalidity. *See Aylus Networks, Inc. v. Apple Inc.*, 856 F.3d 1353, 1362 (Fed. Cir. 2017) ("[S]tatements made by a patent owner during an IPR proceeding, whether before or after an institution decision, can be considered for claim construction and relied upon to support a finding of prosecution disclaimer."); *NFC Tech.*, 2015 WL 1069111, at *7 (determining that even where all claims were not reviewed during IPR proceedings, "any disposition by the PTAB is likely to simplify the proceedings before this Court"). A stay will permit the case to proceed in light of all relevant intrinsic evidence. For example, Orckit will likely make arguments about claim scope to overcome the prior art presented in the IPR petitions, including in its Patent Owner preliminary responses due pre-institution in June 2023. Not staying the case now could mean expending the Court's and the parties' resources on claim construction that could be rendered futile if the PTAB finds certain claims invalid or Orckit takes positions during the IPR proceedings that impact claim scope. *See Anascape, Ltd. V. Microsoft Corp.*, 475 F. Supp. 2d 612, 615 (E.D. Tex. 2007) ("[C]ourts need not expend unnecessary judicial resources by attempting to resolve claims which may be amended, eliminated, or lucidly narrowed by the patent reexamination process and the expertise of its officers."); *Ethicon LLC v. Intuitive Surgical, Inc.*, 2019 WL 1276029, at *2 (D. Del. Mar. 20, 2019) (finding stay pending IPR to be efficient because of additional prosecution history and potential amendments flowing from IPR proceedings).

Finally, granting a stay will potentially simplify the issues for trial through the application of estoppel, which will keep Cisco from pursuing certain invalidity theories after a final written

decision. Even if the PTAB proceedings were to find some of the reviewed claims to be valid, Cisco would still be estopped from asserting invalidity defenses on "any ground that it raised or reasonably could have raised" during *inter partes* review. 35 U.S.C. § 315(e)(2). The exclusion of such art from this litigation provides further simplification warranting a stay here.

D. At Minimum, The Court Should Deny This Motion Without Prejudice So That Cisco Can Renew The Motion After The IPR Institution Decisions.

To the extent the Court finds that the likelihood of simplification is too speculative until the PTAB institutes Cisco's IPRs, the Court should at minimum deny this motion without prejudice so that Cisco can renew its motion. This Court has granted renewed motions to stay following denials without prejudice upon the PTAB later instituting Patent Office challenges. *See Ramot at Tel Aviv Univ. Ltd. v. Cisco Sys., Inc.*, No. 2:19-cv-00225-JRG, 2021 WL 121154, at *1 (E.D. Tex. Jan. 13, 2021) (granting Cisco's renewed motion to stay pending reexamination proceedings after initially denying Cisco's request without prejudice to refile the request if and when any relief by way of the reexams became less speculative or incomplete); *Arbor Glob. Strategies LLC v. Samsung Elecs. Co.*, No. 2:19-cv-00333-JRG-RSP, 2021 WL 66531, at *1 (E.D. Tex. Jan. 7, 2021) (granting renewed motion to stay upon institution of Defendants' IPRs after previously denying motion without prejudice); *e-Watch Inc. v. Apple, Inc.*, No. 2:13-cv-1061-JRG-RSP, 2015 WL 12915668, at *1 (E.D. Tex. Mar. 25, 2015) (granting renewed motion to stay after denying the initial motion to stay "without prejudice to [Defendants'] right to file a motion to stay if the PTAB [] grants the petition to institute"). "The *inter partes* review process [] was designed to give the agency an opportunity to correct its mistakes, to give courts the benefit of the agency's consideration of the effect of prior art on patents being asserted in litigation, and to reduce the burden of litigation on the parties and the courts." *In re Intel Corp.*, No. 2021-168, 2021 WL 4427875, at *2 (Fed. Cir. 2021). For those benefits to apply here, the Court should at minimum

follow its typical practice and "withhold a ruling pending action on the petition by the PTAB or deny the motion without prejudice to refiling in the event that the PTAB institutes a proceeding." *Customedia Techs., LLC*, 2017 WL 3836133, at *1.

V. CONCLUSION

For the foregoing reasons, Cisco respectfully requests that the Court stay this matter pending the PTAB's resolution of the IPRs. However, to the extent the Court follows previous practice, Cisco requests that the Court "withhold a ruling pending action on the petition by the PTAB or deny the motion without prejudice to refiling in the event that the PTAB institutes a proceeding." *Id*.

Dated: March 23, 2023

Respectfully submitted,

/s/ Melissa R. Smith

Melissa R. Smith
GILLAM & SMITH LLP
303 South Washington Avenue
Marshall, TX 75670
Telephone: (903) 934-8450
Facsimile: (903) 934-9257
melissa@gillamsmithlaw.com

Local Counsel for Defendant
CISCO SYSTEMS, INC.

Tamir Packin (*pro hac vice*)
tpackin@desmaraisllp.com
Leslie Spencer (*pro hac vice*)
lspencer@desmaraisllp.com
Deborah J. Mariottini (*pro hac vice*)
dmariottini@desmaraisllp.com
Jordan Owens (*pro hac vice*)
jowens@desmaraisllp.com
DESMARAIS LLP
230 Park Avenue
New York, NY 10169
Tel: (212) 351-3400
Fax: (212) 351-3401

Michael R. Rhodes (*pro hac vice*)
mrhodes@desmaraisllp.com
DESMARAIS LLP
101 California Street
San Francisco, CA 94111
Tel: (415) 573-1900
Fax: (415) 573-1901

Jonathan Lewis (*pro hac vice*)
jlewis@desmaraisllp.com
DESMARAIS LLP
1899 Pennsylvania Avenue, NW
Washington, DC 20006
Tel: (202) 415-4900
Fax: (202) 415-4901

Lead Counsel for Defendant
CISCO SYSTEMS, INC.

CERTIFICATE OF SERVICE

The undersigned hereby certifies that counsel of record who are deemed to have consented to electronic service are being served with a copy of this **DEFENDANT'S MOTION FOR A STAY PENDING ONGOING INTER PARTES REVIEW**, *via* the Court's CM/ECF system per Local Rule CV-5(a)(3) on this the 23rd day of March, 2023.

/s/ Melissa R. Smith

CERTIFICATE OF CONFERENCE

Counsel for Plaintiffs and counsel for Defendants participated in a meet and confer on March 22, 2023, and discussed the relief requested in this motion. Plaintiff is opposed to this motion and therefore the parties are at an impasse requiring Court resolution.

/s/ Melissa R. Smith

EXHIBIT 1

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner

IPR2023-00401
U.S. Patent No. 7,545,740

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 6

I. INTRODUCTION 8

II. GROUNDS FOR STANDING 8

III. NOTE 9

IV. TECHNOLOGY BACKGROUND 9

A. Link Aggregation 9

B. Multiplexing and Demultiplexing 10

V. SUMMARY OF THE '740 PATENT 13

A. Overview of the '740 Patent 13

B. Prosecution History 17

VI. LEVEL OF ORDINARY SKILL IN THE ART 18

VII. CLAIM CONSTRUCTION 18

A. "*interface module*" 19

B. "*selecting, in a single computation*" 19

VIII. RELIEF REQUESTED AND THE REASONS FOR THE REQUESTED RELIEF 20

IX. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE 21

A. Challenged Claims and Statutory Grounds for Challenge 21

B. Grounds 1 & 2: Claims 1-31 are obvious under 35 U.S.C. § 103(a) over Bruckman alone or in view of Basso. 23

1. Summary of Bruckman 23

2. Summary of Basso 27

3. Reasons to Combine Bruckman and Basso 29

4. Summary of Holdsworth 31

5. Holdsworth shows background knowledge of a POSITA 33

6. Claim 1 34

7. Claim 2 47

8. Claim 3 49

9. Claim 4 49

10. Claim 5 50

11. Claim 6 51

12. Claim 7 52

13. Claim 8 53

14. Claim 9 54

15. Claim 10 55

16. Claim 11 58

17. Claim 12 63

18. Claim 13 64

19. Claim 14 66

20. Claim 15 68

21. Claim 16 69

22. Claim 17 70

23. Claim 18 72

24. Claim 19 73

25. Claim 20 73

26. Claims 21-27 73

27. Claim 28 74

28. Claim 29 74

29. Claim 30 75

30. Claim 31 75

C. Grounds 3 & 4: Claims 11 and 26 are obvious under 35 U.S.C. § 103(a) over Bruckman (alone or with Basso) in view of Holdsworth. 76

1. Reasons to Combine Holdsworth with Bruckman 76

2. Claims 11 and 26 79

X. DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE 79

A. Discretionary denial under the *Fintiv* factors is not appropriate 79

1. No evidence regarding a stay 80

2. Parallel proceeding trial date 80

3. Investment in the parallel proceeding 81

4. Overlapping issues with the parallel proceeding 82

5. Identity of parties 82

6. Other circumstances 82

B. Discretionary denial under 35 U.S.C. § 325(d) is not appropriate 83

1. *Becton, Dickinson* Factor (c)..... 84

2. *Becton, Dickinson* Factors (e) and (f): 85

C. Discretionary denial under *General Plastic* is not appropriate..... 87

XI. CONCLUSION.....87

XII. MANDATORY NOTICES 89

A. Real Party-in-Interest 89

B. Related Matters..... 89

C. Lead and Back-up Counsel and Service Information 89

XIII. CLAIMS APPENDIX 91

CERTIFICATE OF WORD COUNT.....105

CERTIFICATE OF SERVICE 106

PETITIONER'S EXHIBIT LIST

Ex.1001	U.S. Patent No. 7,545,740 to Zelig et al.
Ex.1002	Prosecution History of U.S. 7,545,740
Ex.1003	Declaration of Dr. Henry Houh under 37 C.F.R. § 1.68
Ex.1004	*Curriculum Vitae* of Dr. Houh
Ex.1005	U.S. Patent Publication No. 2004/0228278 to Bruckman et al.
Ex.1006	U.S. Patent Publication No. 2003/0210688 to Basso et al.
Ex.1007	Digital Logic Design, by Holdsworth ("Holdsworth") (2002)
Ex.1008	U.S. Patent No. 6,943,580 to Lewis et al. ("Lewis")
Ex.1009	IEEE Standard Dictionary of Electrical and Electronics Terms, (1984)
Ex.1010	RFC 1042, A Standard for the Transmission of IP Datagrams over IEEE 802 Networks
Ex.1011	C++ Inside & Out, by Bruce Eckel ("Eckel") (1993)
Ex.1012	Complaint, *Orckit Corporation v. Cisco Systems, Inc.,* Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Jul. 22, 2022)
Ex.1013	Federal Court Statistics
Ex.1014	Proposed Docket Control Order, *Orckit Corporation v. Cisco Systems, Inc.,* Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Dec. 2, 2022)
Ex.1015	Infringement Contentions, *Orckit Corporation v. Cisco Systems, Inc.,* Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Nov. 3, 2022)
Ex.1016	Affidavit from the Internet Archive
Ex.1017	U.S. Patent No. 7,271,993 to Plojhar ("Plojhar")

Ex.1018	RFC 793, Transmission Control Protocol
Ex.1019	RFC 2026, The Internet Standards Process – Revision 3
Ex.1020	U.S. Patent No. 5,960,428 to Lindsay et al. (“Lindsay”)
Ex.1021	U.S. Patent No. 5,999,538 to Haddock et al. (“Haddock”)
Ex.1022	Course Catalog for Electrical Engineering
Ex.1023	Copyright Record for Eckel (Ex.1011)
Ex.1024	U.S. Patent No. 6,473,424 to DeJager (“DeJager”)

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 1-31 (the "Challenged Claims") of U.S. Patent No. 7,545,740 ("'740 patent," Ex.1001).

The '740 patent relates to link aggregation, which involves joining a group of parallel physical links between two endpoints together into a single logical link. The '740 patent, entitled "Two-way Link Aggregation," was allowed after the Applicant amended the claims to recite that at least one of the links in a link aggregation group was "a bi-directional link operative to communicate in both an upstream direction and a downstream direction." Ex.1002, 59.

As shown below and confirmed in the Declaration of Dr. Houh (Ex.1003), the concept of bi-directional, aggregated links was already known and would have been obvious to a POSITA. *See generally* Ex.1003. The references presented in this Petition render obvious the Challenged Claims, which should be canceled for unpatentability.

II. GROUNDS FOR STANDING

Petitioner certifies that the '740 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.

IV. TECHNOLOGY BACKGROUND

A. Link Aggregation

Link aggregation "is a technique by which a group of parallel physical links between two endpoints in a data network can be joined together into a single logical link." Ex.1005, [0002]. "Link aggregation offers benefits of increased bandwidth, as well as increased availability, since the logical link can continue to function (possibly with reduced bandwidth) even when one of the physical links fails or is taken out of service." Ex.1005, [0002]. Link aggregation for Ethernet networks is defined by industry standards. *See* Ex.1005, [0003].

Link aggregation includes a process for selecting which of the physical links in a Link Aggregation Group (LAG) will transmit a particular frame. This process employs "a distributor function, which distributes data frames submitted by MAC clients among the physical links in the group, and a collector function, which receives frames over the aggregated links and passes them to the appropriate MAC clients." Ex.1005, [0003]. Industry standards describe "possible distribution algorithms that meet the requirements of the standard, while providing some measure of load balancing among the physical links in the aggregation group."

Ex.1005, [0005]. These distribution algorithms commonly use a hash function. *See* Ex.1005, [0006]-[0012].

Hashing techniques used for distributing packets or data frames across link aggregation configurations were well-known before the '740 patent was filed. Ex.1003, ¶¶14-16; *see also* Ex.1024, 5:28.

B. Multiplexing and Demultiplexing

Link aggregation uses a concept referred to as multiplexing. A multiplexer selects from a plurality of input lines for transmission over a single output line. An example 4-input multiplexer implementation is shown below:

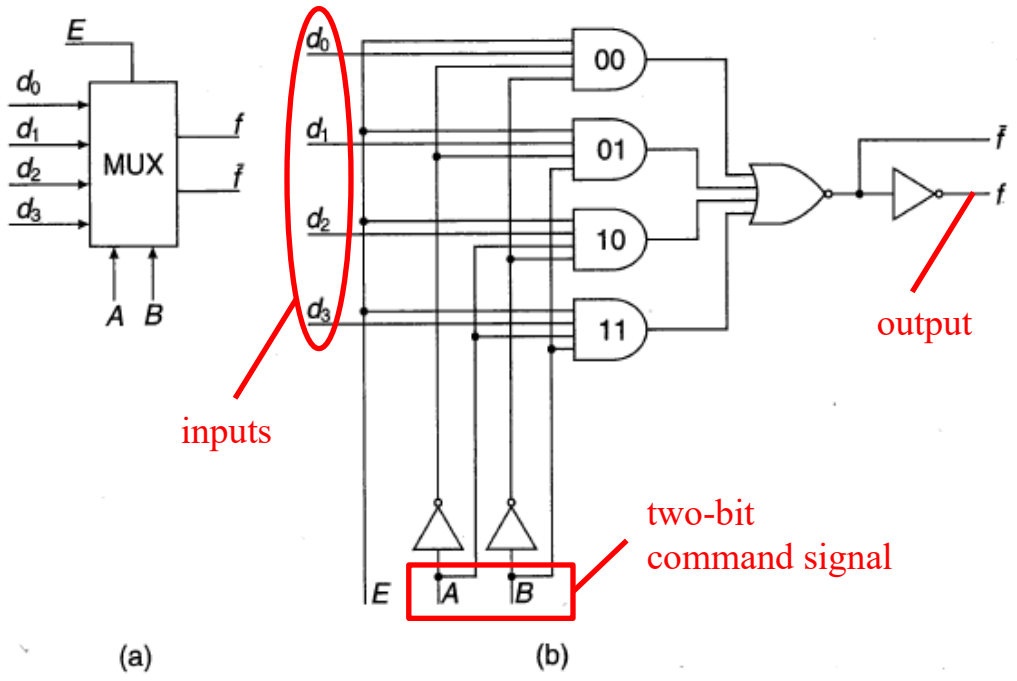


Figure 5.1 *(a) Block diagram of a 4-input multiplexer and (b) its gate implementation*

Ex.1007, Fig. 5.1 (annotated); Ex.1003, ¶17.

A demultiplexer is the opposite of a multiplexer—it selects one of a plurality of output lines for transmission of data received over a single input line. An example 4-output demultiplexer implementation is shown below

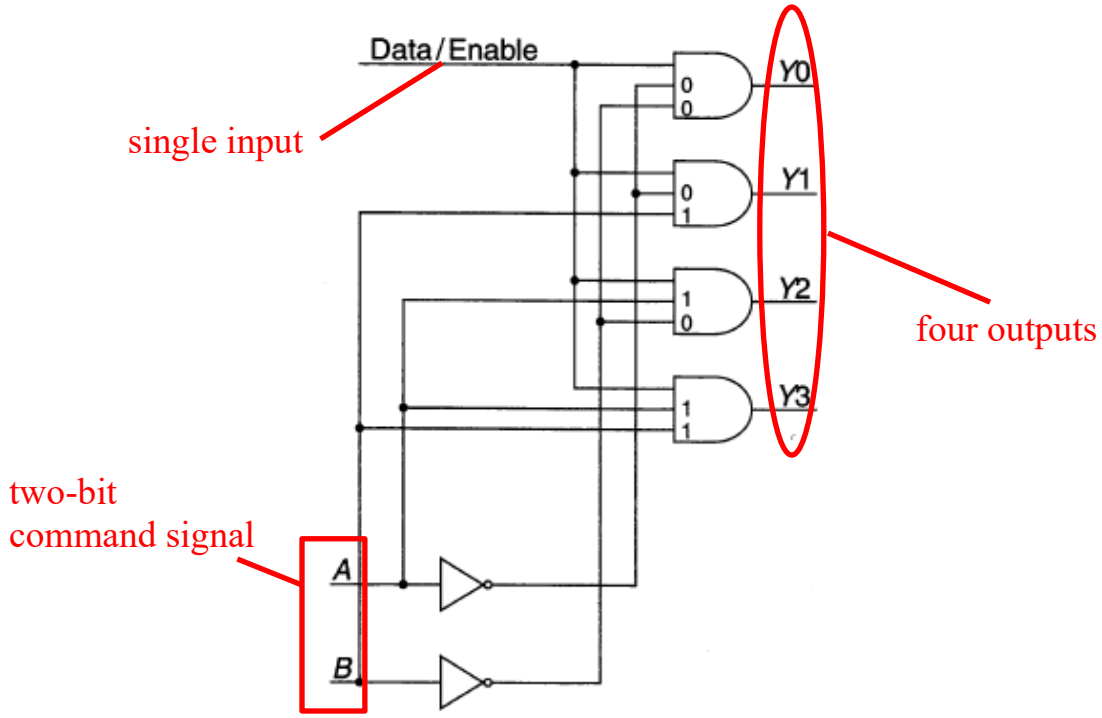
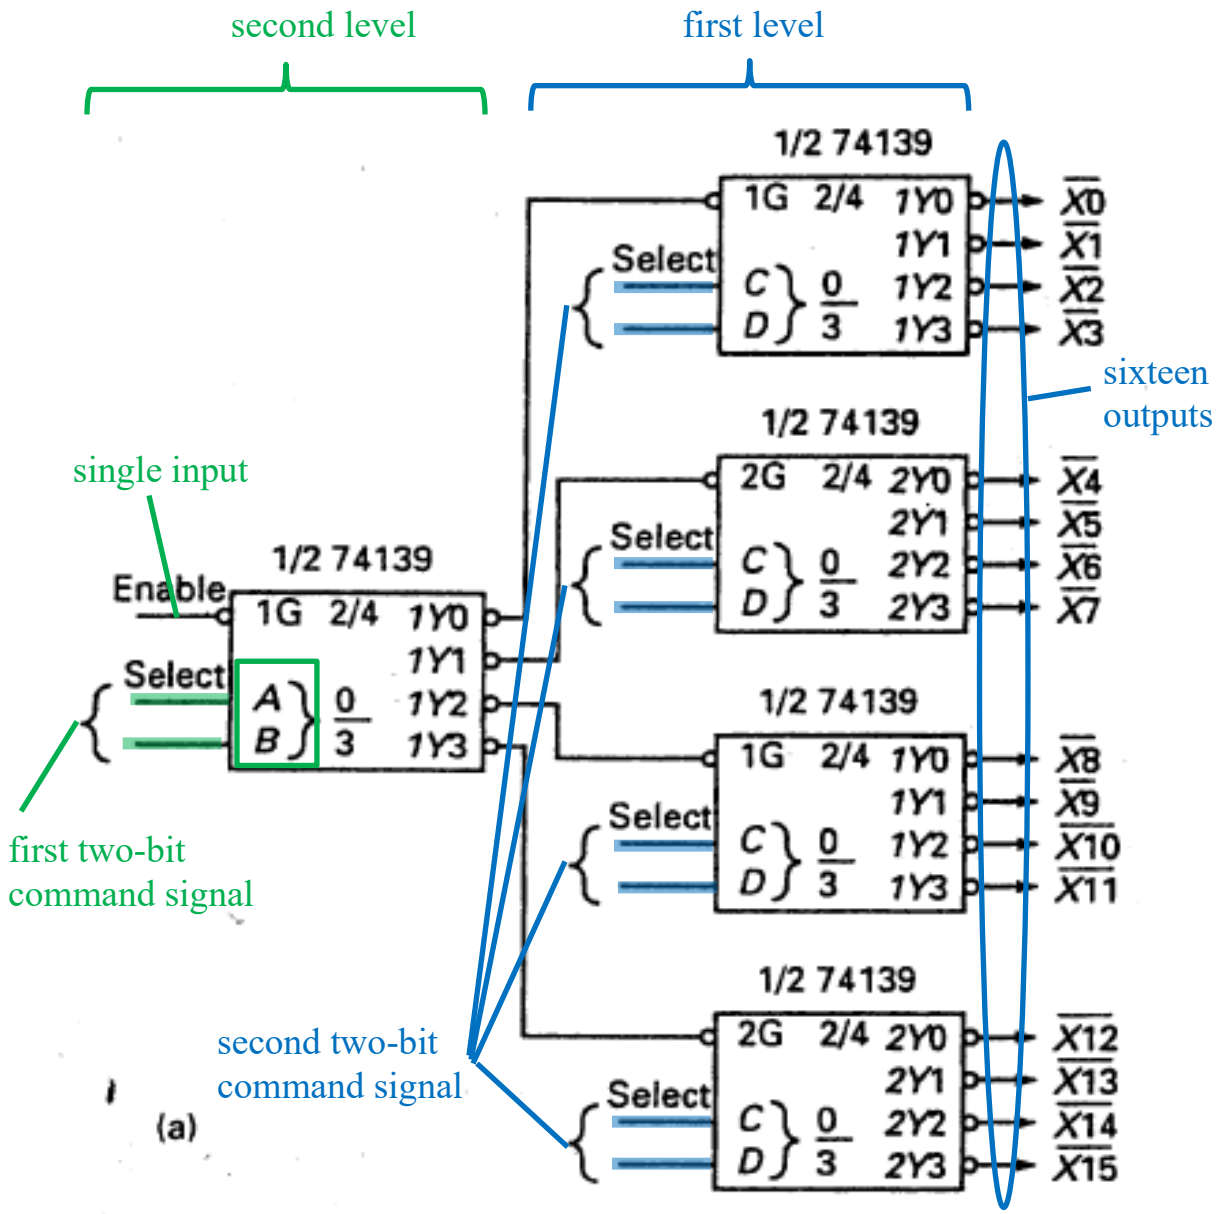


Figure 5.10 *Basic demultiplexer*

Ex.1007, Fig. 5.10 (annotated); Ex.1003, ¶18.

Dr. Houh further explains the functionality of such multiplexers and demultiplexers. Ex.1003, ¶¶17-18.

To form larger de-multiplexers with more output lines, it was known to use two-level structures comprised of several demultiplexers. For example, as shown in the figure below, it was known to use several 1:4 demultiplexers in a two-level structure to create a 1:16 demultiplexer.



Ex.1007, Fig. 5.17(a) (annotated); Ex.1003, ¶19.

As can be seen from the figure above, the first level (on the right) includes four 1:4 demultiplexers that are controlled by lines C and D. The second level (on the left) includes a single 1:4 demultiplexer that is controlled by lines A and B. Through this two-level structure, a four-bit value (A, B, C, D) can select any of the

sixteen possible outputs (X0-X15) for a single input. It was also known to create large multiplexers using a similar two-level structure. *See* Ex.1008, 3:61-7:19; Fig. 8A; Ex.1003, ¶20.

Link aggregation uses demultiplexing to distribute data traffic from a single input among several available outputs. For example, a 1:16 demultiplexer such as the one shown above would be used to distribute data traffic across 16 different parallel output lines. A multiplexer is then used to combine incoming data from all sixteen lines to a single line. Ex.1003, ¶¶17-21.

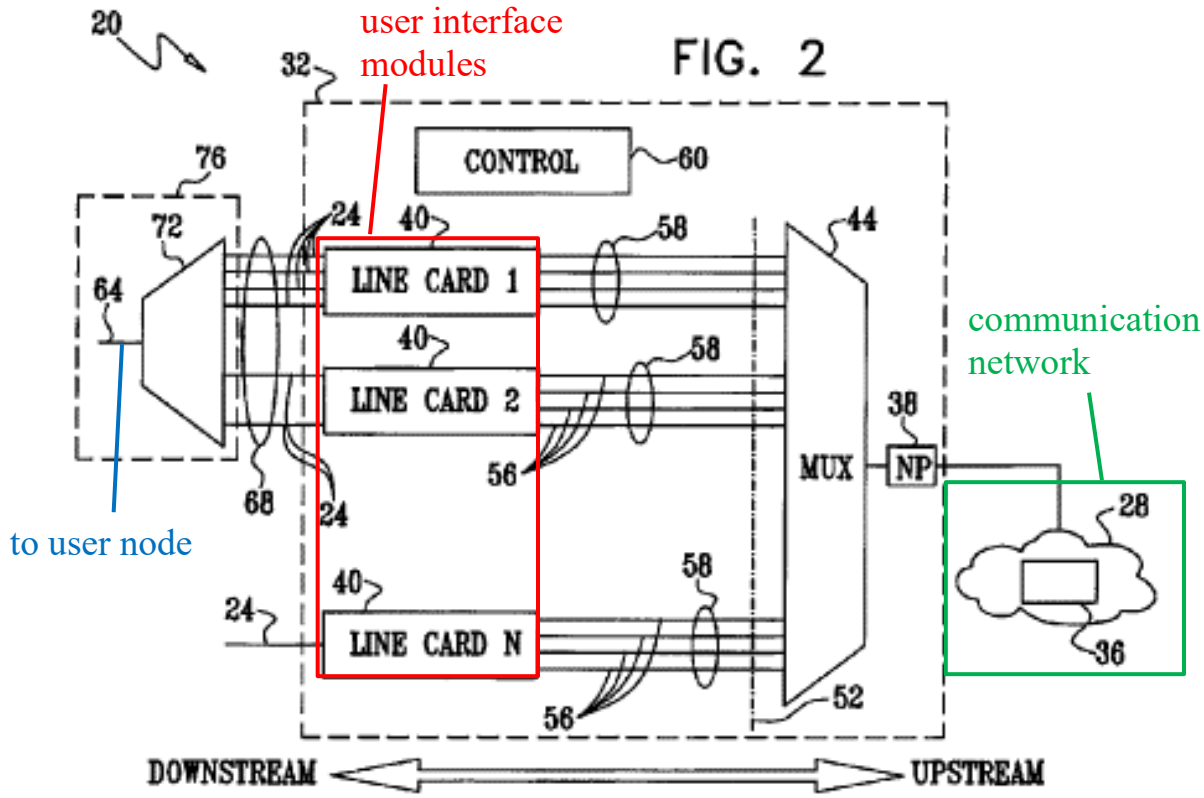
V. SUMMARY OF THE '740 PATENT

A. Overview of the '740 Patent

The '740 patent is "directed to communication networks, and particularly to methods and systems for link aggregation in network elements." Ex.1001, 1:5-7. The '740 patent describes a communication system 20 having a network element 32 that "interconnects a plurality of user ports 24 to a communication network 28." Ex.1001, 4:7-8. "Network 28 may comprise a wide-area network (WAN), such as the Internet, a network internal to a particular organization (Intranet), or any other suitable communication network." Ex.1001, 4:9-11. The user ports may connect "to a user node, such as a layer 2 or layer 3 switch." Ex.1001, 6:38-39.

The network element 32 includes "one or more user interface modules (UIMs), such as line cards 40. Each line card is assigned to process data frames of

one or more user ports." Ex.1001, 4:28-30. Network element 32 is shown below in Fig. 2.

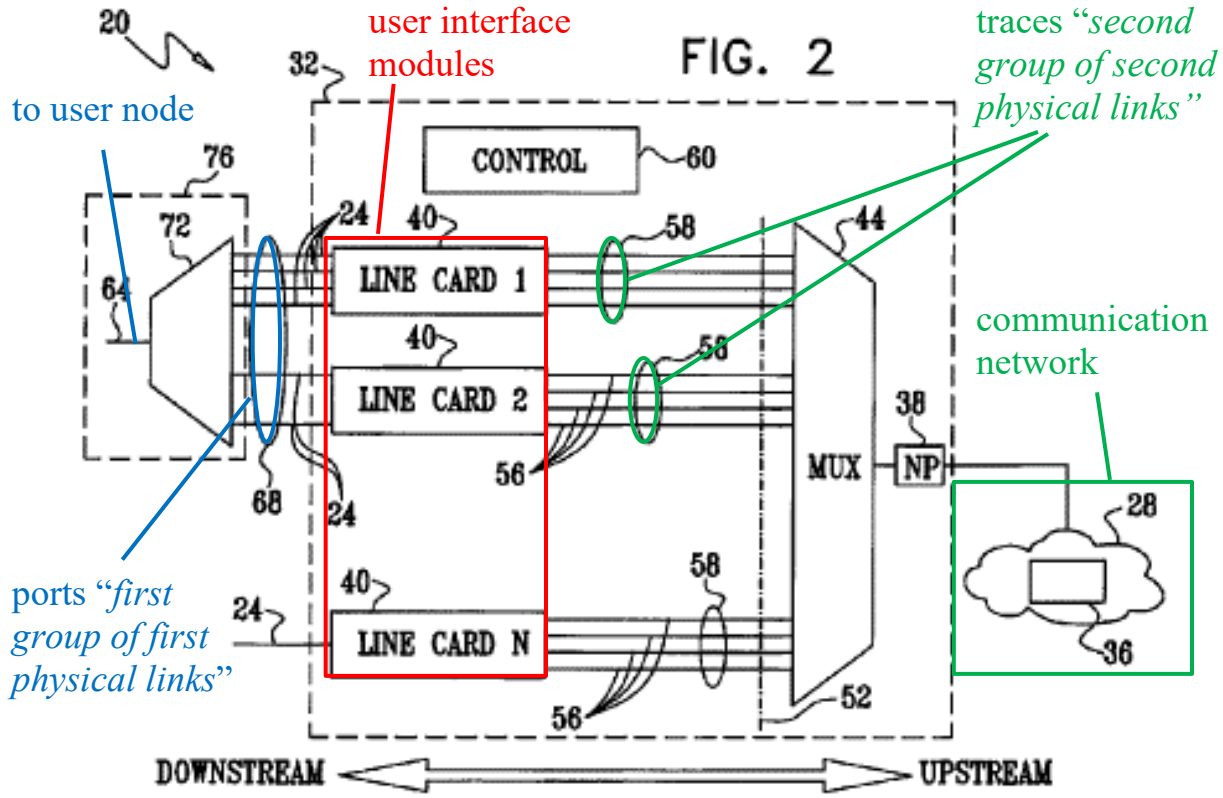


Ex.1001, Fig. 2 (annotated); Ex.1003, ¶23.

The line cards 40 are connected to a user node through a set of ports 24. *See* Ex.1001, 4:5-11, Fig. 2. "User ports 24 forming port 64 are configured as an Ethernet LAG [link aggregation group] group, referred to as an external LAG group 68." Ex.1001, 6:21-23. The '740 patent claims these ports 24 as "*a first group of first physical links arranged in parallel*." *See e.g.,* Ex.1001, claim 1.

The interface modules are connected to the communication network 28 using backplane traces 56: "Backplane 52 comprises physical links, such as

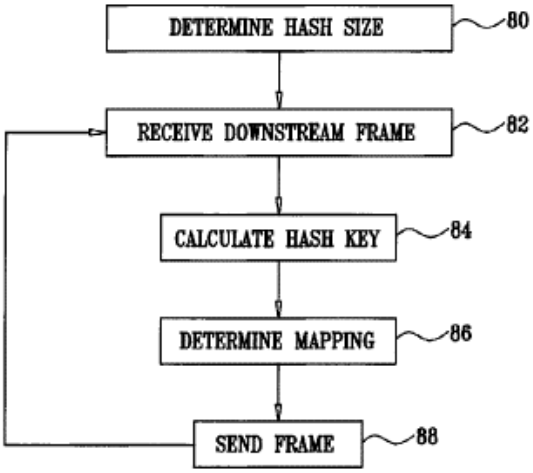
backplane traces 56, typically in the form of printed circuit board (PCB) conductors." Ex.1001, 4:34-37. The '740 patent claims these traces 56 as "*a second group of second physical links arranged in parallel*." *See e.g.,* Ex.1001, claim 1.



Ex.1001, Fig. 2 (annotated); Ex.1003, ¶25.

The network element 32 further includes a control module 60 that, for each received packet, determines which trace and which port will transmit that packet. *See* Ex.1001, 3:45-53. Selection of both a port 24 from the first group and a trace 56 from the second group can be done in "single combined mapping operation that combines the two mapping operations described above." Ex.1001, 6:63-65. The

"combined mapping comprises a single hashing operation that determines, for each such downstream frame, both the backplane trace 56 over which the frame is to be sent to one of line cards 40, and the user port 24 to be used within external LAG group 68." Ex.1001, 6:66-7:3. This combined mapping operation includes several steps, as shown below in Fig. 4.



Ex.1001, Fig. 4.

At step 80, "[t]he method begins with control module 60 determining a hashing size parameter." Ex.1001, 7:25-26. The hashing size parameter is based on the total number of traces multiplied by the total number of ports. *See* Ex.1001, 7:27-30, 42. Then, at step 82, a data frame is received. Ex.1001, 7:43-45. "For each downstream frame, control module 60 calculates a hashing key of the frame… by applying a suitable hashing function to the frame attributes of the downstream frame." Ex.1001, 7:45-50. At step 86, "Module 60 divides the hashing key of the downstream frame by N_{bpow} [the hashing size parameter] and retains the modulo, or

the remainder of the division operation, as a mapping index." Ex.1001, 7:51-53.

Accordingly, the steps shown in Fig. 4 are used to produce a single hash computation that is then used to determine both a port and a trace to transmit the packet. The '740 patent provides an example in which "Module 60 partitions the binary representation of the mapping index into two parts having N1 and N2 bits." Ex.1001, 7:60-61. The first set of bits N1 is used to select one of the ports: "Module 60 uses N1 bits as a user slot/port index, indicating over which user port 24 in external LAG group 68 the frame should be sent." Ex.1001, 7:61-64. The second set of bits N2 is used to select one of the traces: "The remaining N2 bits are used as a backplane trace index, indicating over which of the backplane traces of the relevant line card the frame should be sent." Ex.1001, 7:64-66.

As will be described in further detail below, the link aggregation concepts described and claimed in the '740 patent are rendered obvious by the prior art. Ex.1003, ¶¶22-29.

B. Prosecution History

The '740 patent was filed April 7, 2006. In a first non-final Office Action on October 2, 2008, the Office rejected many claims as anticipated by U.S. Patent No. 6,963,578 to Akahane. Ex.1002, 83. The Office indicated that dependent claims 4, 5, 13, 15, 19, 20, 28, and 30 contained allowable subject matter. Ex.1002, 85.

In response, the Applicant rewrote the dependent claims with allowable

subject matter in independent form. Ex.1002, 76; *see also* Ex.1002, 59-69. Applicant also amended some claims to include limitations requiring the recited communication links to be "bi-directional." Ex.1002, 59-69. The claims were then allowed with no statement of reasons for allowance. Ex.1002, 38.

VI. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill in The Art ("POSITA") in April 2006 would have had a working knowledge of computer networking, including such commonly used technologies as Ethernet and the internet protocol / transmission control protocol (TCP/IP) suite. A POSITA would have had a bachelor's degree in electrical engineering, or an equivalent, and two years of professional experience in the field of communication networks, including link aggregation in communication networks. Lack of professional experience can be remedied by additional education, and vice versa. Ex.1003, ¶¶7-9.

VII. 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). *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). 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). For clarity, however, Petitioner notes below

example embodiments in the specification of certain terms. Ex.1003, ¶¶30-32.

A. "*interface module*"

Each independent claim of the '740 patent recites "*one or more interface modules*." The '740 patent explicitly contemplates that a line card is within the scope of the term *interface module*: "Network element 32 comprises one or more user interface modules (UIMs), such as line cards 40." Ex.1001, 4:28-29. A line card is thus an *interface module* as claimed. Ex.1003, ¶33.

B. "*selecting, in a single computation*"

As explained above at V.A, the '740 patent describes a multi-step selecting process as shown in Fig. 4 and explained in accompanying text. This multi-step process includes determining a hash size (step 80), calculating a hash key (step 84), and determining a mapping (step 86). This multi-step process is recited in various dependent claims—including claims 8-10 which recite that the selecting process includes multiple steps, including (1) "*determining a hashing size responsively to a number of at least some of the first and second physical links*," (2) "*applying the hashing function to the at least one of the frame attributes to produce a hashing key*," and (3) "*calculating a modulo of a division operation of the hashing key by the hashing size*."

It is apparent that the "*single computation*" does not refer to performing the *selecting* process in a single step, but rather to a single computational result of a

hashing function (e.g., single value) that *selects* the outcome. *See, e.g.*, Ex.1001, 5:29-31. Thus, the single value, or the final processing step that produces that single value, are examples of a "*single computation*." Ex.1003, ¶¶34-36.

VIII. 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. Petitioner challenges all claims of the '740 patent because they are all asserted in co-pending litigation. A finding that the Challenged Claims are unpatentable in this proceeding will resolve the parties' dispute in the co-pending litigation and obviate any need for a trial regarding the '740 patent, substantially reducing the time and expense of litigation for all parties.

IX. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

A. Challenged Claims and Statutory Grounds for Challenge[1]

Grounds	Claims	Basis
#1	1-31	35 U.S.C. § 103 (Pre-AIA) over Bruckman
#2	1-31	35 U.S.C. § 103 (Pre-AIA) over Bruckman and Basso
#3	11 and 26	35 U.S.C. § 103 (Pre-AIA) over Bruckman and Holdsworth
#4	11 and 26	35 U.S.C. § 103 (Pre-AIA) over Bruckman, Basso, and Holdsworth

U.S. Patent Publication No. 2004/0228278 to Bruckman et al. ("Bruckman") published on November 18, 2004.

U.S. Patent Publication No. 2003/0210688 to Basso et al. ("Basso") was published on November 13, 2003.

"Digital Logic Design," by Brian Holdsworth and Clive Woods ("Holdsworth") is a textbook with a copyright date of 2002 and is assigned ISBN Number 0-7506-45882. It was published in 2001 and available for sale as of at

[1] For each combination presented herein, Petitioner relies on the teachings, and not on a physical incorporation of elements. *See In re Mouttet*, 686 F.3d 1322, 1332 (Fed. Cir. 2012); *In re Etter*, 756 F.2d 852, 859 (Fed. Cir. 1985).

least August 12, 2001, as evidenced by the Internet Archive—a well-known archiving website. *See* Ex.1016. Holdsworth was also cited in a patent, showing that the book was known to POSITAs. *See* Ex.1017, 6:5-7 ("A useful discussion of logic circuits can be found in "Digital Logic Design", B. Holdsworth and C. Woods, Newnes, 2002"). The totality of the evidence demonstrates that Holdsworth was publicly available to interested persons exercising reasonable diligence, and, therefore, was a printed publication as of 2005. *See Hulu, LLC v. Sound View Innovations, LLC*, IPR2018-01039, Paper 29, at 17-18 (P.T.A.B. Dec. 20, 2018) (precedential) (considering the "totality of the evidence" in assessing whether a reference is prior art). Ex.1003, ¶56.

Bruckman, Basso, and Holdsworth are each prior art under 35 U.S.C. § 102(b).

Petitioner also cites additional prior art as evidence of the background knowledge of a POSITA and to provide contemporaneous context to support Petitioners' 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); *see also K/S HIMPP v. Hear-Wear Techs., LLC*, 751 F.3d 1362, 1365-66 (Fed. Cir. 2014); *Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1363 (Fed. Cir. 2016). For instance, Dr. Houh and this

Petition cite to "Request for Comments" (RFC) documents published by the Internet Engineering Task Force (IETF). To the extent the Board determines that these IETF documents must qualify as prior art "printed publications" for the purposes for which they are cited, the documents do so qualify. *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, ¶38. 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).

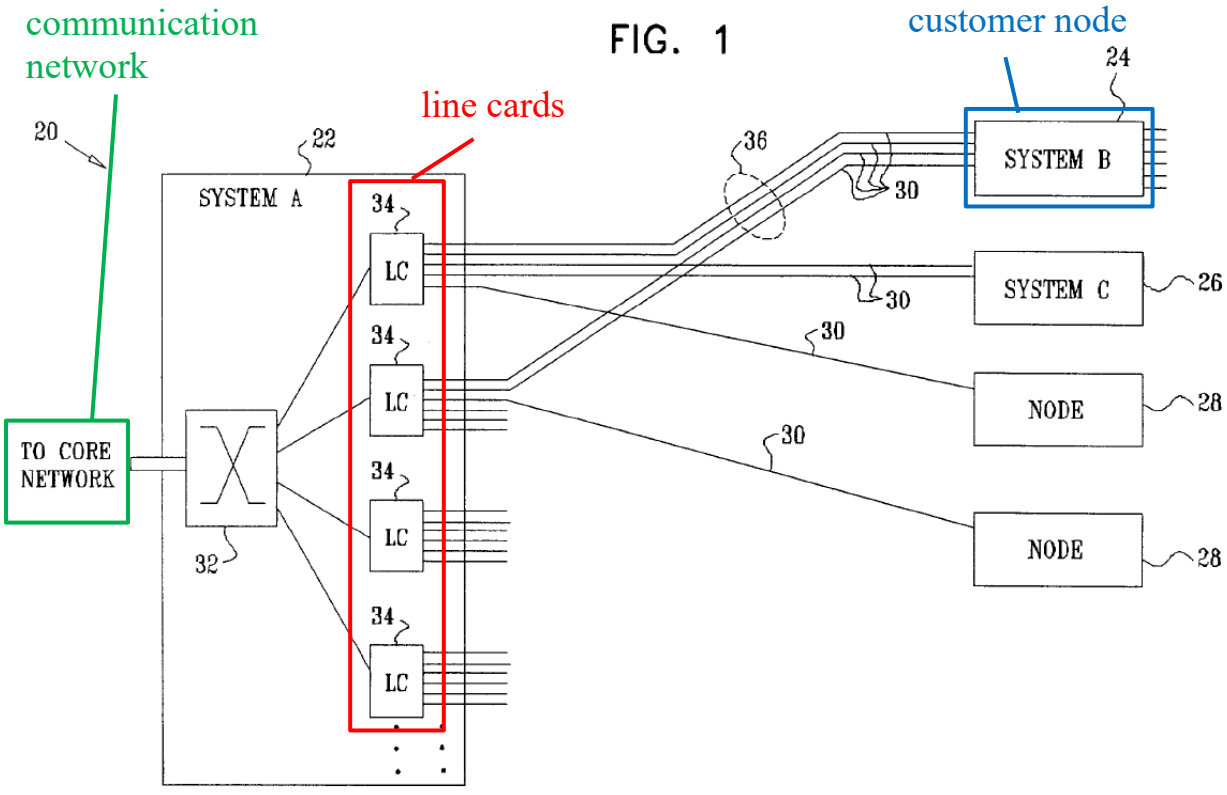
B. Grounds 1 & 2: Claims 1-31 are obvious under 35 U.S.C. § 103(a) over Bruckman alone or in view of Basso.

Analysis for Grounds 1 and 2 is very similar, and thus they are presented together. Ground 2 differs from Ground 1 by relying on additional obviousness teachings from Basso in limitation [1.6]. Ex.1003, ¶¶37-39.

1. Summary of Bruckman

Bruckman relates "to data communication systems, and specifically to methods and systems for link aggregation in a data communication network."

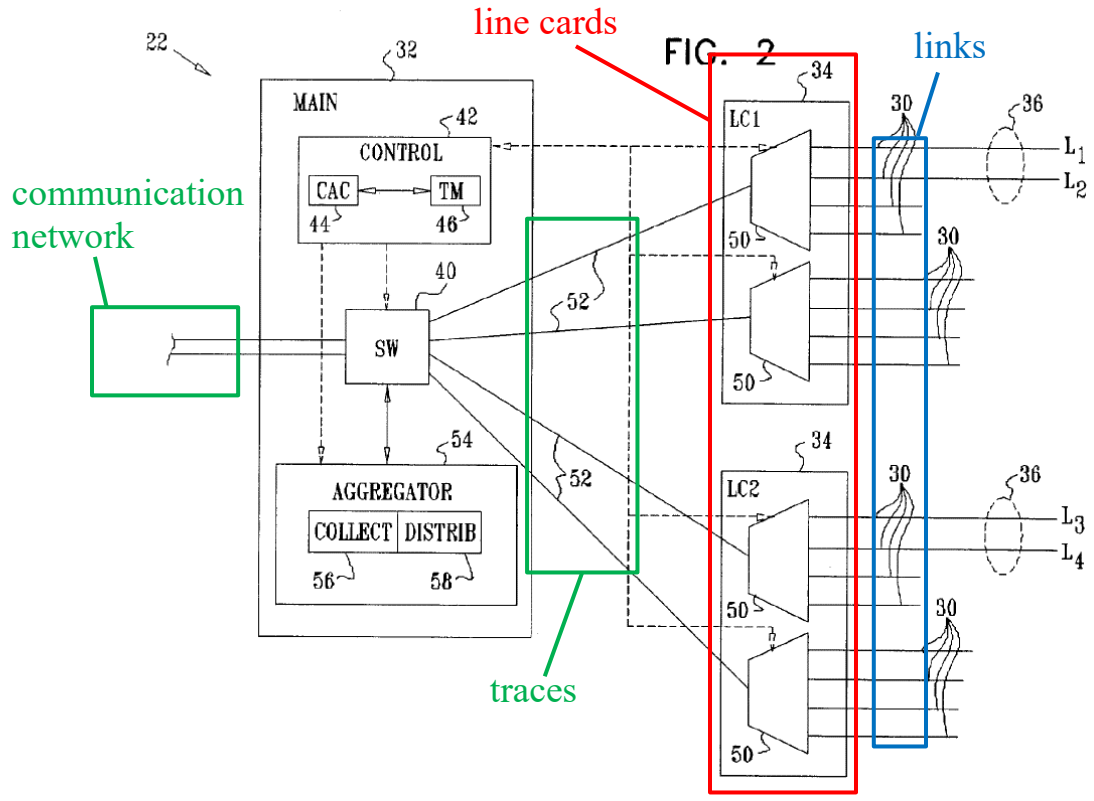
Ex.1005, [0001]. Bruckman describes a communication system 20 having Equipment 22 that "is configured to convey packet data traffic between the customer nodes and a network (which may be a metro network, access network, or other type of core network, for example)." Ex.1005, [0047]. Equipment 22 further includes "a main switching card 32, which is connected to multiple line cards 34." Ex.1005, [0047]. As shown in Fig. 1 below, the main switching card 32 and line cards 34 form a two-level multiplexer/demultiplexer.



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶40.

Fig. 2 (shown below) provides more details of the communication system 20 shown in Fig. 1. The line cards 34 in Fig. 2 connect to the customer node via links

30. "Each line card 34 comprises one or more concentrators 50, which comprise multiple ports that serve respective links 30." Ex.1005, [0056]. The line cards 34 connect to the switching core 40 and core network through backplane traces 52. Bruckman thus describes a two-level multiplexing/demultiplexing structure in which the first level includes the concentrators 50, each of which selects from one of four outputs (links 30) and the second level includes switching core 40, which selects one of a plurality of outputs (traces 52).



Ex.1005, Fig. 2 (annotated); Ex.1003, ¶41.

When processing traffic from the customer nodes to the communication network, Bruckman's two-level structure acts as a multiplexer: "The concentrators

multiplex data traffic between links 30 and traces 52, which connect the concentrators to switching core 40." Ex.1005, [0056]. When processing traffic from the communication network to a customer node, Bruckman's two-level structure acts as a demultiplexer by determining which trace and which port will transmit that data frame. The main card 32 includes a distributor 58, that "determines the link over which to send each frame based on information in the frame header." Ex.1005, [0058]. To make this selection, the distributor "applies a predetermined hash function to the header information." Ex.1005, [0058]. Bruckman provides an example of such a hash function shown in Table 1 below.

```
A  unsigned short hash(unsigned char *hdr, short lagSize)
       {
       short i;
       unsigned short hash=178; // initialization value
       for (i=0; i<4; i++)          // hdr is 4 bytes length
B        hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
C      return (hash % lagSize)
       }
```

Ex.1005, Table 1 (annotated); Ex.1003, ¶42.

Like the selection process in the '740 patent, Bruckman's hashing function includes multiple steps that produce a single hash computation. First, as shown in line A, the hashing function receives the lagSize parameter, which "is the number of active ports (available links 30) in link aggregation group." Ex.1005, [0064]. In

line B, an operation is applied to header information (in variable hdr[2]) to create a hashing key stored in the "hash" variable. Then, in line C, a modulo of the division of the hash key and the LAG size is produced. This final, single modulo value is then used to select which link 30 will transmit the data frame. Ex.1005, [0058].

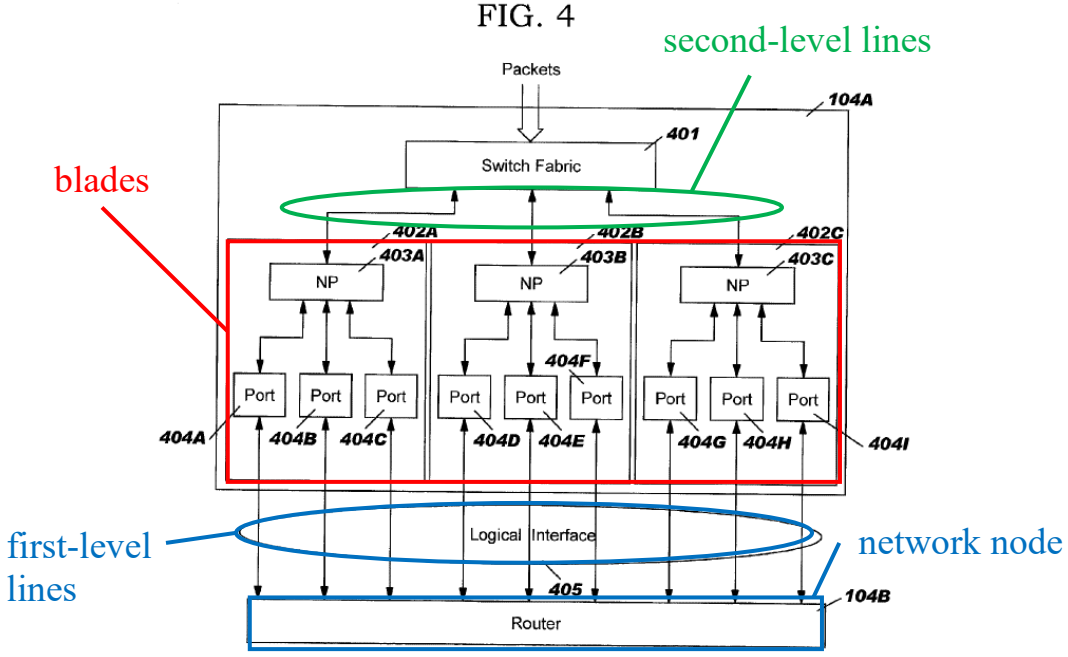
Accordingly, Bruckman shows that using a set of parallel traces and a set of parallel links/ports to transmit packets was known before the filing of the '740 patent. Bruckman additionally shows that it was known to apply a hashing function to select one of the links. Ex.1003, ¶¶40-44.

2. Summary of Basso

Basso relates to link aggregation, particularly, "the field of packet switching networks, and more particularly to logically grouping physical ports of a network device." Ex.1006, [0001]. Basso describes a network device, such as a router, that passes packets between the Internet and a client device. Like Bruckman, Basso describes a two-level structure, as shown in Fig. 1 below.

[2] "Here hdr is the header of the frame to be distributed." Ex.1005, [0064].

FIG. 4



Ex.1006, Fig. 4 (annotated); Ex.1003, ¶45.

Basso's network device 104a includes "a switch fabric coupled to a plurality of blades where **each blade may comprise** one or more network processors coupled to **one or more ports**." Ex.1006, [0009]. As shown in Fig. 4, Basso's two-level structure includes a first group of physical links—each associated with one of the ports 404—that connect the network device to a network node (in this case, Router 104b). Additionally, a second set of links connects each of the blades to the switching fabric: "[A] switch fabric 401 configured to direct the incoming packets of data to particular **blades 402A-C coupled to switch fabric 401**." Ex.1006, [0030].

Upon receipt of a packet, "a hash function may be performed on the source

and destination address in the packet header to generate a hash value." Ex.1006, [0011]. Using that single hash value, "an appropriate **blade/port combination** may be identified to transmit the received packet of data." Basso, [0012]. Basso thus uses a single hash computation to select both a first-level line (one of the ports) and second-level line (one of the blades)—a blade/port combination. Ex.1003, ¶¶45-47.

3. Reasons to Combine Bruckman and Basso

A POSITA would have found it obvious that when using a two-level demultiplexer structure such as the one described by Bruckman, a single hash computation would be used to control *both* a first-level and a second-level demultiplexer. Basso provides explicit evidence of doing so. Ex.1003, ¶48.

As an initial matter, both Bruckman and Basso are analogous art to the '740 patent. Bruckman and Basso are directed to the same field of endeavor as the '740 patent—network communications. *Compare* Ex.1001, 1:5-7 ("The present invention relates generally to communication networks, and particularly to methods and systems for link aggregation in network elements") *with* Ex.1005, [0001] ("The present invention relates generally to data communication systems, and specifically to methods and systems for link aggregation in a data communication network.") *and* Ex.1006, [0001] ("The present invention relates to the field of packet switching networks, and more particularly to logically grouping physical ports of a network device, e.g., router, into logical interfaces."); *see also*

Ex.1003, ¶48.

Bruckman's hashing technique produces a single hash computation that is used to determine which link 30 receives a data frame: The "distributor 58, which is responsible for distributing data frames arriving from the network **among links 30** in aggregation group 36 … determines **the link** over which to send each frame based on information in the frame header." Ex.1005, [0058].

Bruckman provides limited description of how the hash computation is used to select a trace. Basso provides further details regarding such two-level selections and includes explicit evidence of using a single hash computation to make a selection in both levels of a two-level multiplexer/demultiplexer structure. As explained above, Basso's network device uses one hash value to select "an appropriate **blade/port combination** … to transmit the received packet of data." Ex.1006, [0012]. Basso thus uses a single hash computation to select a blade/port combination—both a first-level line (one of the ports) and second-level line (one of the blades). Accordingly, to the extent a POSITA implementing the device of Bruckman needed more information about how a single hash computation would be used to control both levels of a two-level multiplexer/demultiplexer structure, such teaching is provided by Basso.

A POSITA would have had a reasonable expectation of success because Basso describes a substantially similar two-level multiplexer/demultiplexer

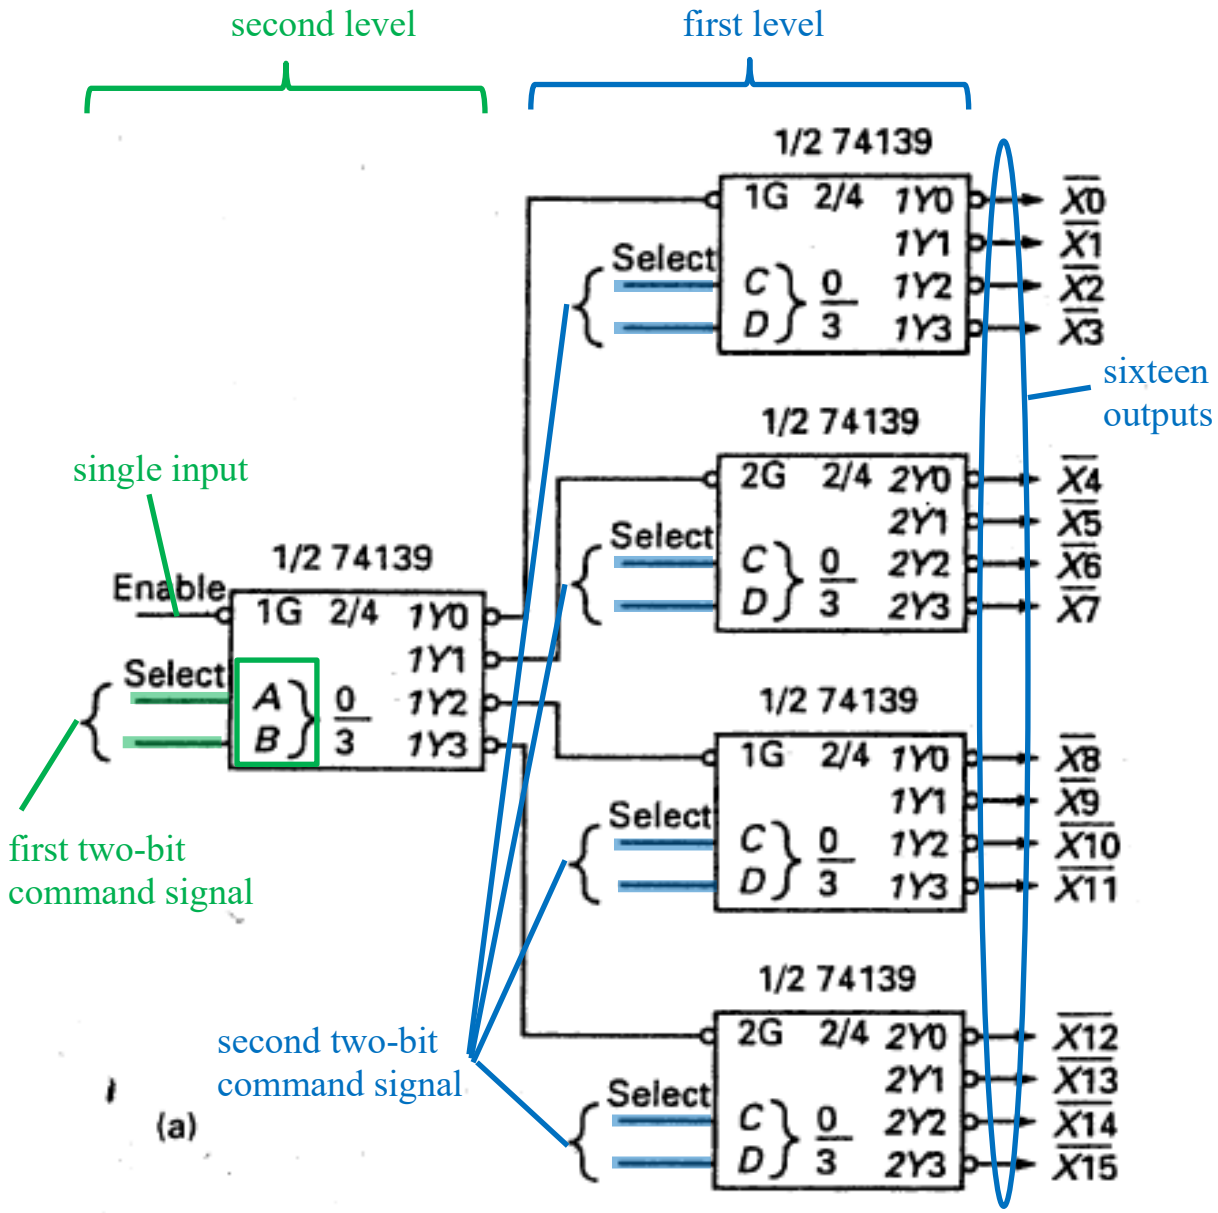
structure as Bruckman. Ex.1003, ¶52. Thus, using a single hash computation to make selections in both levels of the two-level multiplexer/demultiplexer structure would have worked as well in Bruckman. A POSITA would have further found the combination beneficial as using a single hash computation (rather than multiple) would have allowed for improved operational efficiency due to the simplified design. Ex.1003, ¶52. Calculating a single hash value as opposed to multiple hash values reduces the computational load and thus allows for faster switching, which is consistent with Bruckman's goal to "ensure that sufficient bandwidth will be available on the links in the group in order to meet service guarantees." Ex.1005, [0015].

Thus, the combination of Bruckman and Basso represents the application of a known technique (Basso's single hash computation to select both levels in a two-level multiplexer/demultiplexer structure) to Bruckman's method (using a single hash computation for a two-level multiplexer/demultiplexer structure) to yield predictable results (fast and computationally efficient switching). Ex.1003, ¶¶48-53.

4. Summary of Holdsworth

Holdsworth is a textbook entitled "Digital Logic Design." Holdsworth describes a basic multiplexer and a basic demultiplexer. *See* Ex.1007, Fig. 5.1 and 5.10. Holdsworth further describes how 1:4 demultiplexers may be used in a two-

level fashion to create a 1:16 multiplexer, as shown in Fig. 5.17(a) below.



Ex.1007, Fig. 5.17(a) (annotated); Ex.1003, ¶54.

Holdsworth thus shows that when using two levels of 1:4 demultiplexers to create a 1:16 demultiplexer, a designer would use a first two-bit signal for selection at the first level and a second two-bit signal for selection at the second level. In

other words, Holdsworth shows that it was a well-known technique to use a first subset of bits for selection at the first level of a two-level demultiplexer, and a second subset of bits for selection at the second level of the two-level demultiplexer.

Holdsworth is analogous art to the '740 patent, as its engineering principles are applicable to the same field of endeavor as the '740 patent—communication networks. As explained above at IV.B, link aggregation in communication networks utilizes multiplexers, which Holdsworth describes in detail. *See generally,* Ex.1007. Because Holdsworth provides implementation details for multiplexers and demultiplexers, Holdsworth is reasonably pertinent to the technical problem of load balancing and link aggregation to which the '740 patent relates. Ex.1003, ¶¶54-56.

5. Holdsworth shows background knowledge of a POSITA

The Holdsworth textbook is "intended to cover all the material that is needed in a typical undergraduate or master's course on Digital Logic Systems, and also to act as a reference text for graduates working in this field." Ex.1007, preface. Holdsworth represents the background knowledge of a POSITA because Holdsworth's subject matter would have been part of the undergraduate education of a POSITA. Ex.1003, ¶57. For example, digital logic systems was a commonly required course in undergraduate electrical engineering degree plans. *See* Ex.1022.

6. Claim 1

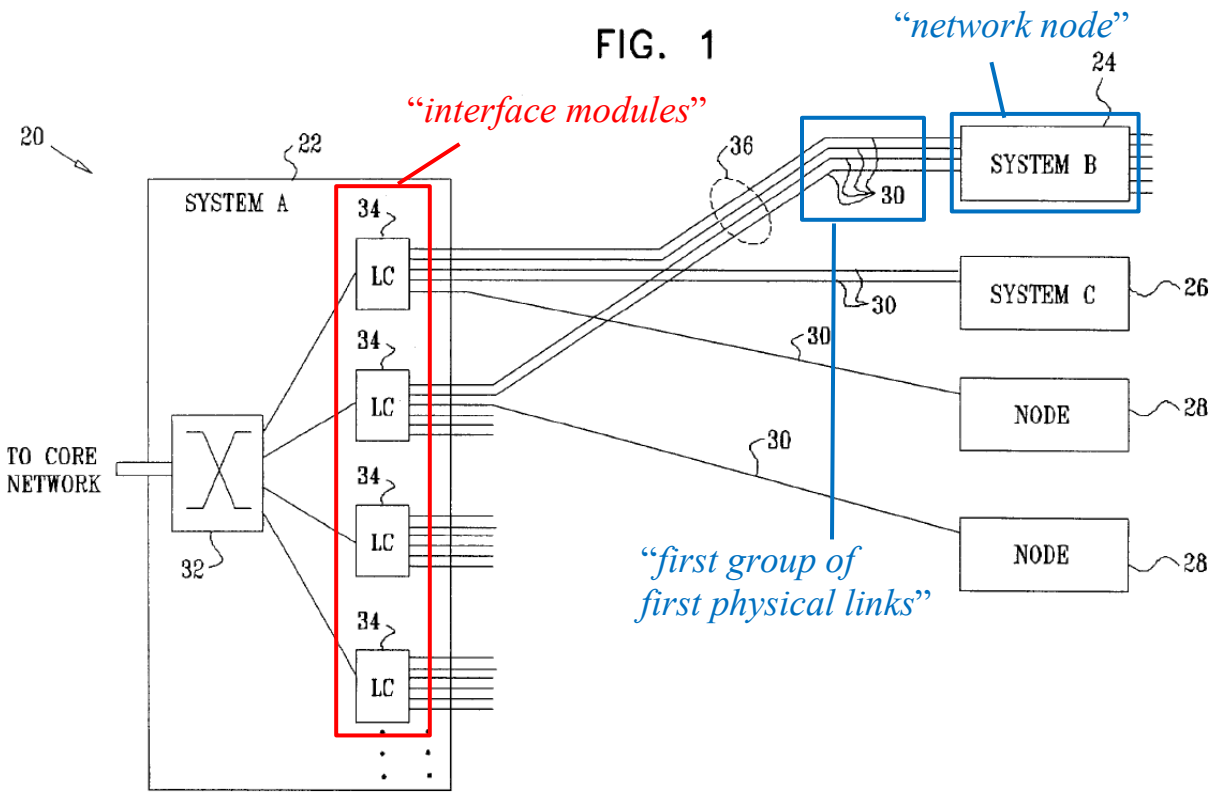
[1.0] ***A method for communication, comprising:***

Bruckman "relates generally to data communication systems, and specifically to **methods** and systems for link aggregation in a **data communication** network." Ex.1005, [0001]; abstract.

Thus, because Bruckman describes methods for link aggregation in a data communication network, Bruckman renders obvious a "*method for communication*" as claimed. Ex.1003, ¶¶58-59.

[1.1] ***coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel,***

Bruckman describes a communication system that includes a set of networked customer nodes 24, 26, 28 (any of which is a "*network node*") that are connected to multiple line cards 34 ("*one or more interface modules*") over parallel physical links 30 ("*a first group of physical links arranged in parallel*"). "In this example, central office equipment 22 communicates with customer nodes 24, 26, 28,... over physical links 30… For this purpose, equipment 22 comprises a main switching card 32, which is connected to multiple line cards 34 that serve links 30." Ex.1005, [0047].



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶60.

The links 30 shown in Fig. 1 form an "aggregation group 36." Bruckman, [0057]. "Link aggregation is a technique by which a group of **parallel physical links** between two endpoints in a data network can be joined together into a single logical link." Ex.1005, [0002].

Thus, because Bruckman describes connecting a customer node to a set of line cards over a set of parallel links, Bruckman renders obvious "*coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel*" as claimed. Ex.1003, ¶¶60-62.

[1.2] ***at least one of said first physical links being a bi-directional link operative***

to communicate in both an upstream direction and a downstream direction;

Bruckman's Equipment 22 transmits data over the physical links 30 in both directions. "Links 30 typically comprise **full-duplex** Ethernet links." Ex.1005, [0047]. The term full-duplex "[r]efers to a communication system or equipment capable of transmission simultaneously in two directions." Ex.1009, 376; *see also* Ex.1021, 8:11-13. Accordingly, the links 30 transmit data in both directions.

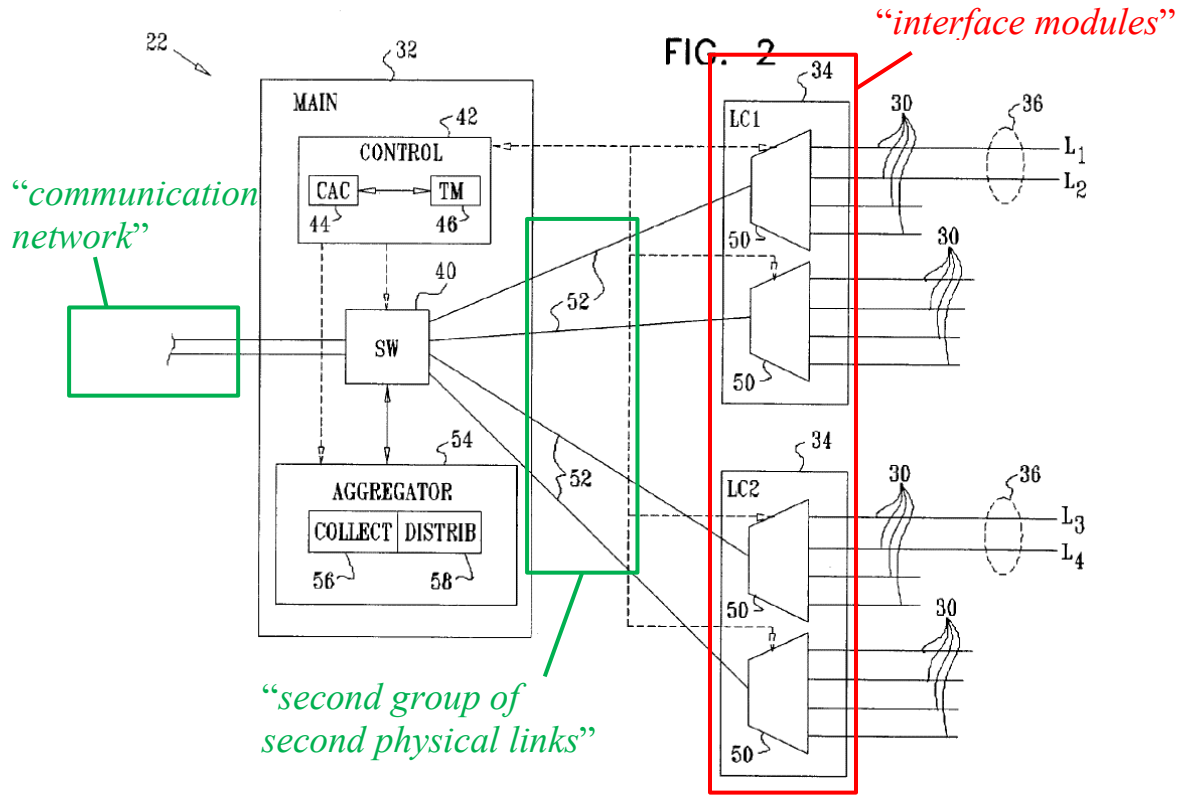
For data traffic in the downstream direction (from communication network to customer node), Bruckman's equipment 22 uses a distributor 58: "Aggregator 54 comprises **a distributor 58, which is responsible for distributing data frames arriving from the network among links 30** in aggregation group 36." Ex.1005, [0058]. In the upstream direction (from customer node to communication network), Bruckman utilizes a collector 56: "Aggregator 54 further comprises a collector 56, **which collects data frames that were received over different links 30** in group 36, and arranges the frames back into a single traffic stream." Ex.1005 [0065].

Thus, because Bruckman describes traffic flowing in both directions across full-duplex links, Bruckman renders obvious "*at least one of said first physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction*" as claimed. Ex.1003, ¶¶63-66.

[1.3] ***coupling each of the one or more interface modules to a communication***

network using a second group of second physical links arranged in parallel,

In Bruckman's system, the line cards 34 ("*one or more interface modules*") are coupled to the core network ("*communication network*") with a set of parallel traces 52 ("*a second group of second physical links arranged in parallel*"). As shown in Fig. 2 below, "Each line card 34 comprises one or more concentrators 50, which comprise multiple ports that serve respective links 30. The concentrators multiplex data traffic between links 30 and traces 52, which connect the concentrators to switching core 40." Ex.1005, [0056].

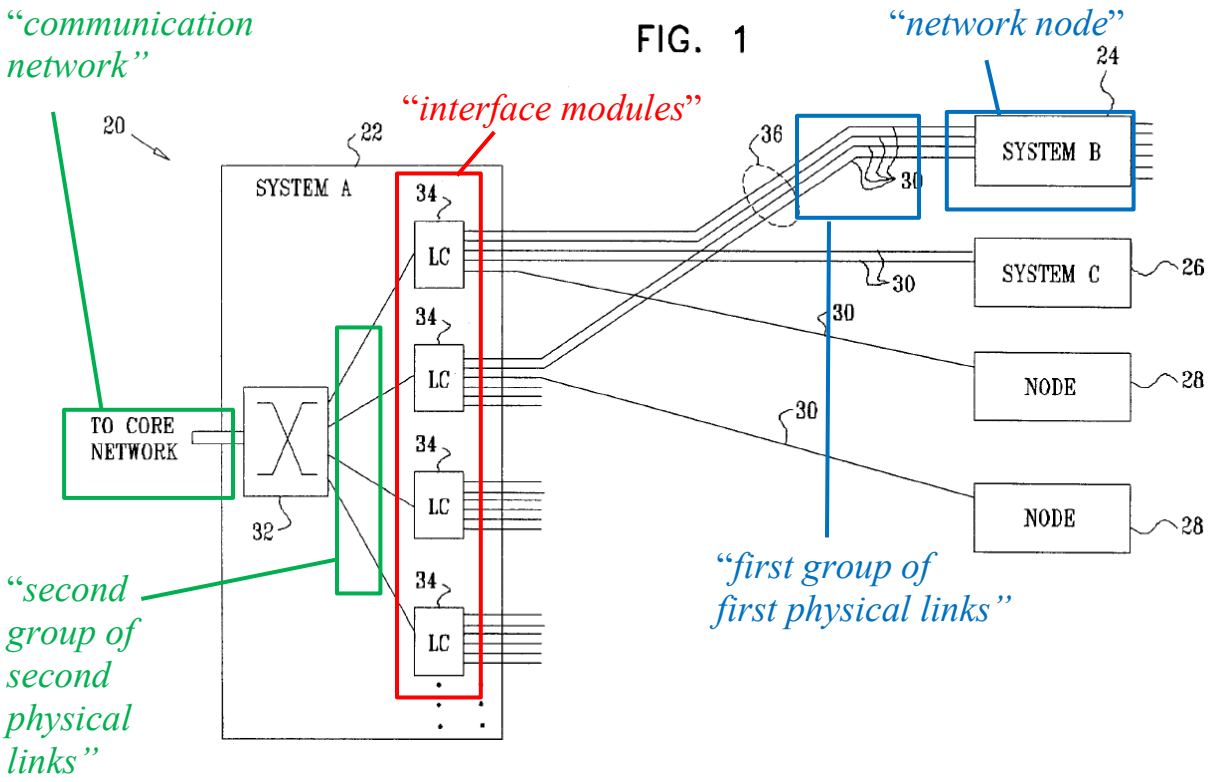


Ex.1005, Fig. 2 (annotated); Ex.1003, ¶67.

Each of the traces 52 connects the switching core 40 to one of the

concentrators 50 and are arranged in parallel in Fig. 2. Traces 52 are part of a link aggregation group because Bruckman describes how "aggregation group extends over a number of concentrators 50," and the traces 52 couple the concentrators to the switching core 40. Ex.1005, [0066]; *see also id.* [0057]. Since the traces 52 are part of a link aggregation group, Bruckman further renders obvious that the traces are "a group of parallel physical links." Ex.1005, [0002].

As shown in Fig. 1 below, the main card 32 connects to the core network. Thus, by connecting the line cards 34 ("*interface modules*") to the switching core 40 of the main card 32, the traces 52 connect the line cards 34 to a "*communication network*" as claimed.



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶69.

Thus, because Bruckman's line cards are connected to a core network using a set of parallel traces, Bruckman renders obvious "*coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel*" as claimed. Ex.1003, ¶¶67-70.

[1.4] *at least one of said second physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction;*

As discussed at [1.2], Bruckman's system handles traffic in both directions between the core network and the customer nodes. An aggregator distributes traffic not only to full-duplex links 30 (*see* [1.2]), "but also to traces 52 that connect to

multiplexers 50 that serve these links." Ex.1005, [0057]. Since the aggregator includes both a traffic "distributor" and a traffic "collector," it would have been obvious that traffic flows in both directions over traces 52. Ex.1005, [0058], [0065]; Ex.1003, ¶71.

Accordingly, at least one of the traces 52 ("*at least one of said second physical links*") is "*a bi-directional link operative to communicate in both an upstream direction and a downstream direction*" as claimed. Ex.1003, ¶72.

[1.5] ***receiving a data frame having frame attributes sent between the communication network and the network node;***

Bruckman's main card 32 includes an aggregator 54 that receives data frames that are sent between the core network ("*communication network*") and links 30 connected to customer nodes (a "*network node*"). The data frames include various attributes, including frame headers ("*frame attributes*"): "Typically, distributor 58 determines the link over which to send each frame based on information in the **frame header**, as described in the Background of the Invention. Ex.1005, [0058].

Thus, because Bruckman's main card 32 receives data frames from the core network for transmission to the links 30 which connect to customer nodes, and those data frames have headers, Bruckman renders obvious "*receiving a data frame having frame attributes sent between the communication network and the*

network node" as claimed. Ex.1003, ¶¶73-74.

[1.6] ***selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and***

First, Bruckman's aggregator 54 includes a distributor 58 that "determines the link over which to send each frame." Ex.1005, [0058]. "Preferably, distributor 58 applies a predetermined hash function to the header information ['*frame attributes*']." Ex.1005, [0058].

Bruckman gives an example of a hash function to select a link in Table 1, below:

```
unsigned short hash(unsigned char *hdr, short lagSize)
  {
  short i;
  unsigned short hash=178; // initialization value
  for (i=0; i<4; i++)          // hdr is 4 bytes length
    hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
  return (hash % lagSize)
  }
```

"*single computation*"

Ex.1005, Table 1, [0063] (annotated); Ex.1003, ¶76.

Bruckman's hashing function thus selects which link 30 ("*first physical link out of the first group*") will transmit the data frame from the core network to the customer node based on the value calculated from the "hash % lagSize" operation ("*in a single computation*"). A POSITA would have understood that the final line

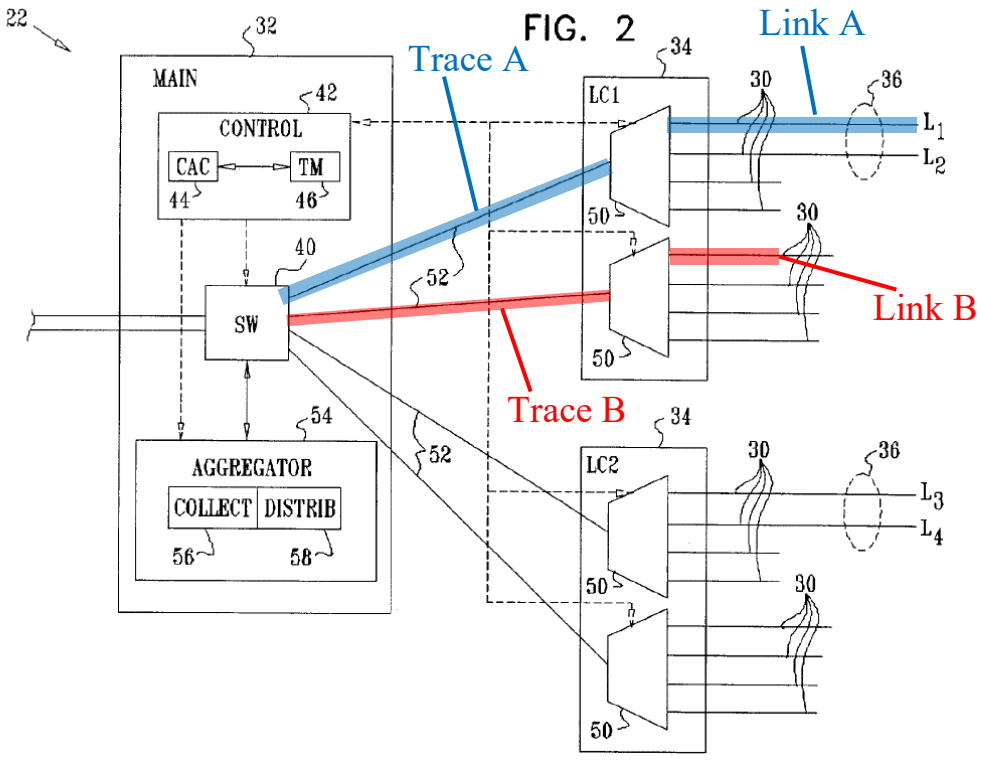
of code in this hashing function returns the value of the operation "hash % lagSize." Ex.1003, ¶77. A POSITA would have recognized that the code snippet provided in Table 1 is written using the syntax of the C or C++ programming languages. In both C and C++ languages, the % operator refers to a mathematical modulus operation, which returns the remainder from integer division. *See* Ex.1011[3], 85; Ex.1020, 7:44-45; Ex.1003, ¶77. Either the modulus operation "hash % lagSize" itself or its result render obvious the claimed "*single computation*."

Second, a POSITA would have found it obvious that Bruckman's hashing function selects both which link 30 ("*first physical link out of the first group*") **and** which trace 52 ("*second physical link out of the second group*") will transmit the data frame from the core network to the customer node. Ex.1003, ¶78. Bruckman explicitly states that the hash function determines which link 30 receives the data frame: The "distributor 58, which is responsible for distributing data frames arriving from the network **among links 30** in aggregation group 36 … determines **the link** over which to send each frame based on information in the frame header." Bruckman, [0058].

Bruckman provides limited description of how the hash computation is used

[3] Ex.1011 has a publication date of 12-18-1992 according to the copyright record. Ex.1023.

to select a trace. A POSITA would have recognized, however, that Bruckman's selection of a link 30 also determines which trace 52 will transmit the data frame. Ex.1003, ¶79. For example, referring to Fig. 2 below, if Link A is selected by the hash function, then Trace A is also selected since Trace A is the only trace that connects to link A to the switching core. Similarly, if Link B is selected by the hash function, then Trace B is also selected, since Trace B is the only trace that connects link B to the switching core. Accordingly, selection of a specific link 30 also determines a specific trace 52, e.g., whether to select Trace A or Trace B.



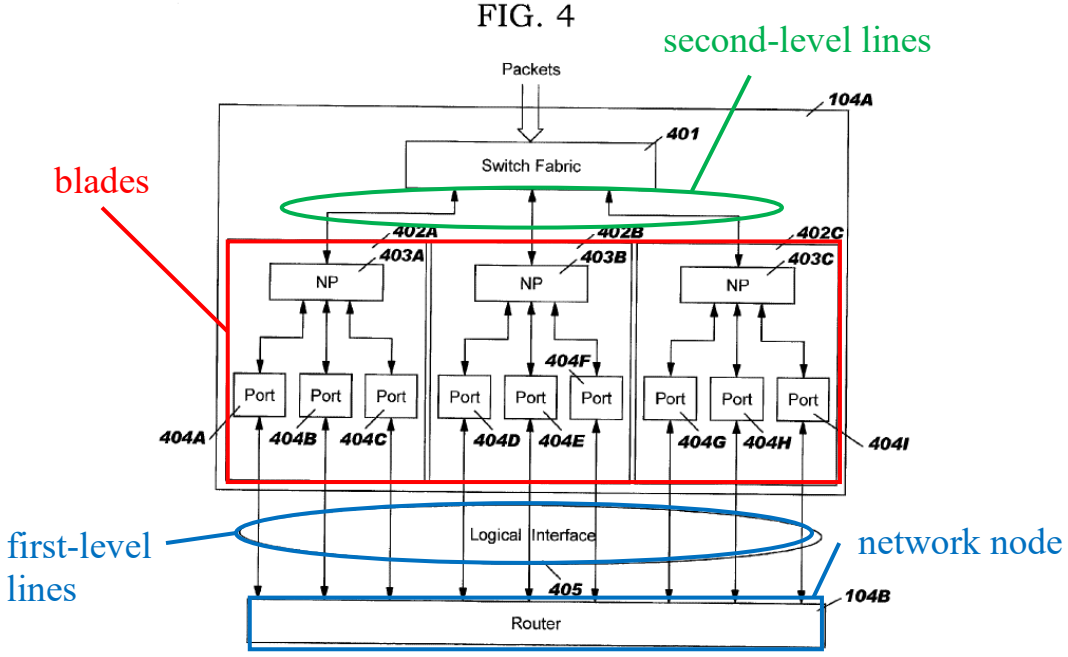
Ex.1005, Fig. 2 (annotated); Ex.1003, ¶79.

A POSITA would have found it obvious that selection of a link 30 also

determines which trace 52 will transmit the data frame because each link 30 is associated with only a single trace 52. Ex.1003, ¶80.

Ground 2: Alternatively, it would have been obvious to use the result of a single hash computation to select links from two groups from the combination of Bruckman and Basso. Basso describes a two-level demultiplexer that is analogous to Bruckman's two-level demultiplexer, as shown in Fig. 4 below.

FIG. 4



Ex.1006, Fig. 4 (annotated); Ex.1003, ¶82.

Like Bruckman's links 30, Basso describes a set of physical links between blades of a router 104a to another network node (router 104B in the example of Fig. 4). Like Bruckman's traces 52, each of Basso's blades is connected to the switch fabric, which is connected to a communication network such as the Internet.

See Ex.1006, Fig. 1.

Basso describes using a single hash computation to select a blade/port combination. Upon receipt of a packet, "a hash function may be performed on the source and destination address in the packet header to generate a hash value." Ex.1006, [0011]. Using that single hash value, Basso's network device selects "an appropriate **blade/port combination** ... identified to transmit the received packet of data." Ex.1006, [0012]. Basso thus uses a single hash computation to select both a first-level line (one of the ports) and second-level line (one of the blades)—a blade/port combination.

Selection of Bruckman's blade/port combination is analogous to selection of a trace 30 and link 52 in Bruckman. For the reasons explained above at IX.B.3, a POSITA would have found it obvious for Bruckman's hash computation to select a trace and link combination, as evidenced by Basso's hash computation that similarly selects a blade/port combination.

Thus, Bruckman alone or together with Basso renders obvious "*selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group*" as claimed. Ex.1003, ¶¶75-84.

[1.7] ***sending the data frame over the selected first and second physical links,***

Bruckman explains that its distributor "is responsible for **distributing data**

frames arriving from the network among links 30 in aggregation group 36." Ex.1005, [0058].

Consistent with the example above in [1.6], it would have been obvious that sending a frame over a given link (such as Link B) involves sending the data frame over the single trace (such as Trace B) that leads to that link as well. *See supra*, [1.6]; Ex.1003, ¶86.

Thus, because Bruckman describes transmitting frames over a specific link based on the result of the hash function, Bruckman renders obvious "*sending the data frame over the selected first and second physical links*" as claimed. Ex.1003, ¶¶85-87.

[1.8] ***said sending comprising communicating along at least one of said bi-directional links.***

As explained above at [1.2] and [1.4], Bruckman's device transmits data frames in both directions, and thus it would have been obvious for links 30 and traces 52 to be bi-directional. A POSITA would have thus understood that traffic passed from the Bruckman's core network to the customer node is communicated across bi-directional links. Thus, Bruckman renders obvious "*said sending comprising communicating along at least one of said bi-directional links*" as claimed. Ex.1003, ¶88.

7. Claim 2

[2.1] ***The method according to claim 1, wherein the network node comprises a user node, and***

Bruckman's customer nodes 24, 26, and 28 are "*user node[s]*" as claimed. Indeed, Bruckman uses the term customer and user interchangeably. *Compare* Ex.1005, [0013] ("**Service level agreements** between network service providers and **customers** commonly specify a certain committed bandwidth") *with* Ex.1005, [0072] ("The connections on links 30, including any link aggregation groups, then compete for the remaining available bandwidth (typically in a weighted manner, based on the amount of excess bandwidth contracted for in the **users' service level agreements**, as is known in the art)").

Thus, because Bruckman's nodes 24, 26, and 28 are referred to as customer nodes, and Bruckman uses customer and user interchangeably, Bruckman renders obvious "*wherein the network node comprises a user node*" as claimed. Ex.1003, ¶¶89-90.

[2.2] ***wherein sending the data frame comprises establishing a communication service between the user node and the communication network.***

First, as explained above at [1.7], Bruckman's device transmits data frames to customer nodes over the links selected from the link aggregation group.

Second, Bruckman's disclosure is directed to "**establishing** a connection with a guaranteed bandwidth for transmitting data over a logical link that includes

a plurality of parallel physical links between first and second endpoints." Ex.1005, abstract; *see also* Ex.1005, claims 1, 11, 19, 32. Bruckman explains that establishment of the aggregated links "provides the total bandwidth guaranteed by the customer's service level agreement" and provides "for bandwidth allocation in a link aggregation system to ensure that sufficient bandwidth will be available on the links in the group in order to meet service guarantees."
Ex.1005, [0014]-[0015].

A POSITA would have understood that the presence of an agreement means that there would have been a service established before the frames are actually sent and that sending the frames is in furtherance of that established service. Ex.1003, ¶93. Alternatively, a POSITA would have found it obvious for the data traffic sent via Bruckman's system to include packets used to establish a common network communication service, a transmission control protocol (TCP) session. Ex.1003, ¶93 (citing Ex.1018). As of the priority date, TCP communications were one of the most common type of network communication service in use. Ex.1003, ¶93.

Thus, Bruckman's communication of data frames across the aggregated links provides a service between the communication network and customer nodes, which renders obvious "*wherein sending the data frame comprises establishing a communication service between the user node and the communication network*" as claimed. Ex.1003, ¶¶91-94.

8. Claim 3

[3.1] ***The method according to claim 1, wherein the second physical links comprise backplane traces formed on a backplane to which the one or more interface modules are coupled.***

First, Bruckman's traces 52 ("*second physical links*") are part of a "printed circuit back plane" and are thus "*backplane traces formed on a backplane*" as claimed. "Typically, main card 32 and line cards 34 are arranged in a card rack and plug into a printed circuit back plane, (not shown) which comprises traces 52." Ex.1005, [0056].

Second, because the line cards 34 ("*one or more interface modules*") are "arranged in a card rack and plug into [the] printed circuit back plane," the line cards are "*coupled*" to the printed circuit back plane. Ex.1005, [0056].

Thus, Bruckman renders obvious "*wherein the second physical links comprise backplane traces formed on a backplane to which the one or more interface modules are coupled*" as claimed. Ex.1003, ¶¶95-97.

9. Claim 4

[4.0]-[4.5] ***A method for … second physical links,***

See [1.0]-[1.7]. Ex.1003, ¶¶98-103.

[4.6] ***at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.***

First, Bruckman explains that the links 30 ("*first group of physical links*") may be organized into link aggregation groups. "For example, an aggregation

group 36 of four physical links is defined between equipment 22 and node 24." Ex.1005, [0048]. "In the example shown in FIG. 2, aggregation group 36 comprises links L1 and L2, which are connected to LC1, and links L3 and L4, which are connected to LC2." Ex.1005, [0057].

Second, Bruckman explains that "the embodiments described herein refer specifically to link aggregation in Ethernet (IEEE 802.3) networks." Ex.1005, [0017]; *see also* [0003] ("For Ethernet networks, link aggregation is defined by Clause 43 of IEEE Standard 802.3").

Thus, because Bruckman's set of links 30 are Ethernet links that form a link aggregation group, Bruckman renders obvious "*at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group*" as claimed. Ex.1003, ¶¶104-06.

10. Claim 5

[5.0]-[5.5] ***A method for … and second physical links,***

See [1.0]-[1.7]. Ex.1003, ¶¶107-112.

[5.6] ***coupling the network node to the one or more interface modules comprises aggregating two or more of the first physical links into an external Ethernet link aggregation (LAG) group***

See [4.6]. Ex.1003, ¶113.

[5.7] ***so as to increase a data bandwidth provided to the network node.***

Bruckman explains that "[l]ink aggregation offers benefits of increased

bandwidth." Ex.1005, [0002].

Thus, because link aggregation is used to increase bandwidth, Bruckman renders obvious using a LAG "*so as to increase a data bandwidth provided to the network node*" as claimed. Ex.1003, ¶¶114-15.

11. Claim 6

[6.1] ***The method according to claim 1, wherein coupling each of the one or more interface modules to the communication network comprises at least one of multiplexing upstream data frames sent from the network node to the communication network, and***

Bruckman's line cards 34 ("*interface modules*") include concentrators that "multiplex data traffic between links 30 and traces 52." Ex.1005, [0056]. Bruckman also uses the term concentrators and multiplexers interchangeably. *See* Ex.1005, [0056] (referring to "concentrators 50"); Ex.1005, [0057] (referring to "multiplexers 50").

Thus, because Bruckman's concentrators/multiplexers 50 multiplex traffic from the customer nodes to the communication network, Bruckman renders obvious "*wherein coupling each of the one or more interface modules to the communication network comprises at least one of multiplexing upstream data frames sent from the network node to the communication network*" as claimed. Ex.1003, ¶¶116-18.

[6.2] ***[at least one of…] demultiplexing downstream data frames sent from the communication network to the network node.***

A POSITA would have understood that demultiplexing is the opposite of multiplexing. *See* Ex.1003, ¶119. As explained above at [1.2] and [1.4], Bruckman's links and traces are bi-directional. *See* Ex.1005, [0058]. A POSITA would have understood that when the line card 34 receives a signal on a particular trace, and then selects one of several links with which to transmit the signal, the line card 34 is performing a demultiplexing function. Ex.1003, ¶119. In other words, while traffic in the upstream direction (customer node to communication network) is multiplexed by the line cards 34, downstream traffic (communication network to customer node) is demultiplexed by the line cards 34.

Thus, because Bruckman's links are bi-directional and traffic in the downstream direction is demultiplexed by the line cards, Bruckman renders obvious "*demultiplexing downstream data frames sent from the communication network to the network node*" as claimed. Ex.1003, ¶120.

12. Claim 7

[7.1] ***The method according to claim 1, wherein selecting the first and second physical links comprises balancing a frame data rate among at least some of the first and second physical links.***

Bruckman's hashing technique for selecting a link "distributes traffic in an approximately uniform manner across the entire set of possible hash values." Ex.1005, [0061]. By distributing the load equally among possible hash values (which correspond to specific links), Bruckman's hashing technique is designed to

balance data among the links. Moreover, Bruckman incorporates by reference known algorithms for load balancing data traffic among links. *See* Ex.1005, [0005].

Thus, because Bruckman's hashing technique seeks to distribute the load across the links in a uniform manner, Bruckman renders obvious "*wherein selecting the first and second physical links comprises balancing a frame data rate among at least some of the first and second physical links*" as claimed. Ex.1003, ¶¶121-22.

13. Claim 8

[8.1] ***The method according to claim 1, wherein selecting the first and second physical links comprises applying a mapping function to the at least one of the frame attributes.***

As explained above at [1.6], Bruckman's selection process includes applying a hashing function to the header. Bruckman's hashing function is shown in Table 1 below:

hashing function "*mapping function*"

DISTRIBUTOR HASH FUNCTION

```
unsigned short hash(unsigned char *hdr, short lagSize)
  {
  short i;
  unsigned short hash=178; // initialization value
  for (i=0; i<4; i++)          // hdr is 4 bytes length
    hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
  return (hash % lagSize)
  }
```

Ex.1005, Table 1 (annotated); Ex.1003, ¶123.

Bruckman's hashing function is a mapping function because it maps a particular data frame to a specific trace 52 and link 30 through the system. Further, claim 9 explicitly recites that a hashing function is within the scope of a "*mapping function*."

Thus, because Bruckman's selecting process involves a hashing function, Bruckman renders obvious "*wherein selecting the first and second physical links comprises applying a mapping function to the at least one of the frame attributes*" as claimed. Ex.1003, ¶¶123-25.

14. Claim 9

[9.1] *The method according to claim 8, wherein applying the mapping function comprises applying a hashing function.*

As explained above at [8.1], Bruckman's hash function is a mapping function. Thus, Bruckman renders obvious "*wherein applying the mapping function comprises applying a hashing function*" as claimed. Ex.1003, ¶¶126-27.

15. Claim 10

[10.1] ***The method according to claim 9, wherein applying the hashing function comprises determining a hashing size responsively to a number of at least some of the first and second physical links,***

As shown in line A, Bruckman's hash function receives two parameters: hdr and lagSize. Bruckman explains that "lagsize is the number of active ports (available links 30) in link aggregation group." Ex.1005, [0064].

DISTRIBUTOR HASH FUNCTION

```
A   unsigned short hash(unsigned char *hdr, short lagSize)
    {
    short i;
    unsigned short hash=178; // initialization value
    for (i=0; i<4; i++)          // hdr is 4 bytes length
      hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
    return (hash % lagSize)
    }
```

"*determining a hashing size*"

Ex.1005, Table 1 (annotated); Ex.1003, ¶128.

Accordingly, by receiving the lagSize parameter, which is the number of active ports (available links 30) in the link aggregation group, Bruckman's hashing technique determines "*a hashing size responsively to a number of at least some of the first and second physical links*." Thus, Bruckman renders this limitation obvious. Ex.1003, ¶129.

[10.2] ***applying the hashing function to the at least one of the frame attributes to produce a hashing key,***

As shown on line A of the figure below, Bruckman's hashing function receives the "hdr" parameter ("*frame attributes*"). "Here hdr is the header of the frame to be distributed." Bruckman, [0064]. Then, in Line B, the hdr value is used to produce a "*hashing key*" represented by the variable "hash."

DISTRIBUTOR HASH FUNCTION

```
A   unsigned short hash(unsigned char *hdr, short lagSize)
      {
      short i;
      unsigned short hash=178; // initialization value
      for (i=0; i<4; i++)          // hdr is 4 bytes length
B       hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
      return (hash % lagSize)
      }
```

variable representing "*frame attributes*"

"*hashing key*"

Ex.1005, Table 1 (annotated); Ex.1003, ¶130.

The "hash" variable represents a *hashing key* because it is used in a modulus operation (%) to produce the final hash result. *See supra* [1.6] & *infra* [10.3]. A POSITA would have been familiar with programming terminology and recognized that in many programming languages, including the code shown in Table 1, the % symbol represents the modulus operation. *See* Ex.1011, 85; Ex.1003, ¶131.

Thus, because Bruckman's hashing function uses the header information to produce a hashing key, Bruckman renders obvious "*applying the hashing function to the at least one of the frame attributes to produce a hashing key*" as claimed. Ex..1003, ¶132.

[10.3] ***calculating a modulo of a division operation of the hashing key by the hashing size, and***

As shown in the annotated Bruckman hashing algorithm below, Line C includes the operation "hash % lagSize." As known to POSITAs, the % operator represents the modulus (remainder) operation used to calculate the remainder from integer division. *See* Ex.1011, 85; Ex.1020, 7:44-45; *supra* [1.6]. A POSITA would have thus understood that the operation in Line C calculates the remainder ("*modulo*") that results from dividing the value of "hash" ("*hashing key*") by the value of "lagSize" ("*hashing size*"). Ex.1003, ¶133.

DISTRIBUTOR HASH FUNCTION

```
A   unsigned short hash(unsigned char *hdr, short lagSize)
      {
      short i;
      unsigned short hash=178; // initialization value
      for (i=0; i<4; i++)          // hdr is 4 bytes length
B       hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
C     return (hash % lagSize);
      }
```

"*hashing size*"

"*hashing key*"

Bruckman, Table 1.

Bruckman's hashing function thus calculates and returns a "*modulo of a division operation*." Where there are 16 different links (as in Bruckman's example), the value of the lagSize parameter ("*hashing size*") would be 16. Ex.1003, ¶134.

Thus, Bruckman's hashing function, which returns "hash % lagSize," renders obvious "*calculating a modulo of a division operation of the hashing key by the hashing size*" as claimed. Ex.1003, ¶135.

[10.4] ***selecting the first and second physical links responsively to the modulo.***

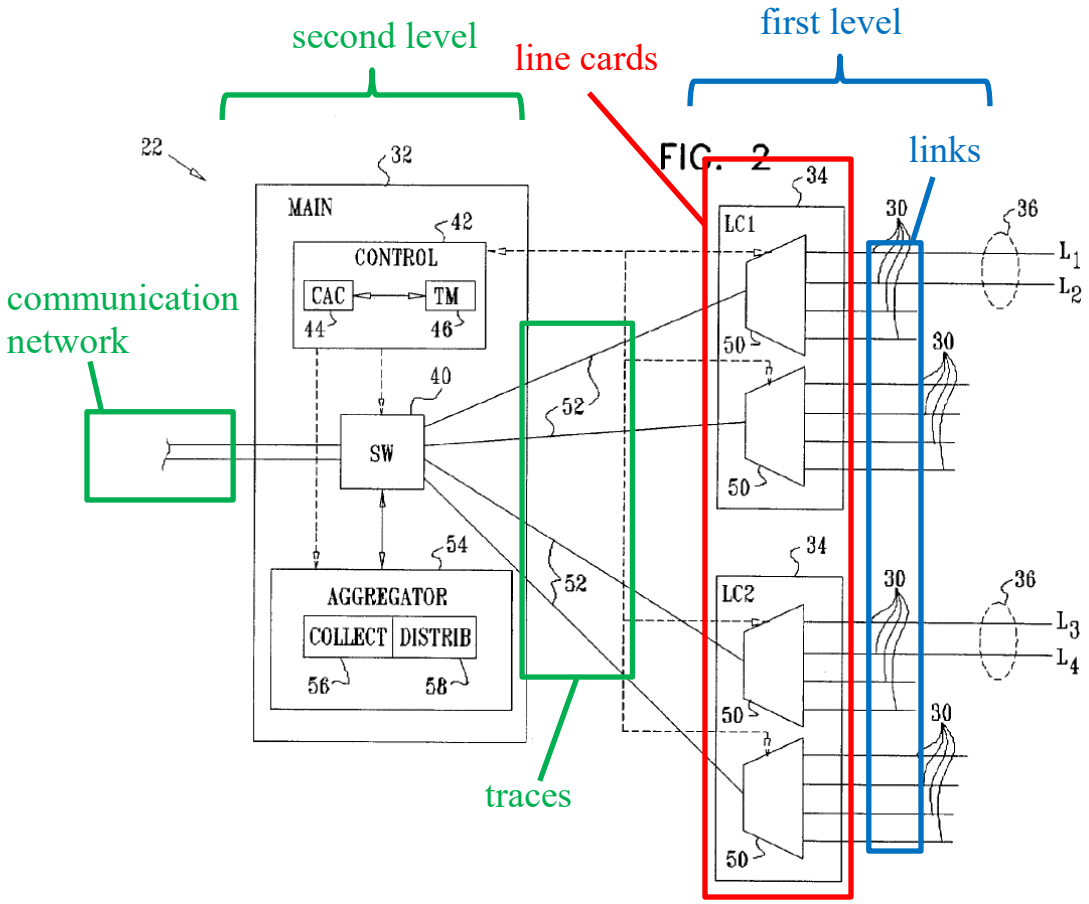
For the reasons explained above at [1.6], Bruckman's application of a hashing function teaches selecting first and second physical links responsive to the single hash result. Bruckman's hashing function returns a modulo, as shown in [10.3] by the line "return (hash % lagSize)." The modulus operation, or its resulting value (the modulo), is the single computation that is used to select a trace and a link as described above at [1.6]. Thus, Bruckman renders obvious "*selecting the first and second physical links responsively to the modulo*" as claimed. Ex.1003, ¶136.

16. Claim 11

[11.1] ***The method according to claim 10, wherein selecting the first and second physical links responsively to the modulo comprises selecting the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo.***

This claim limitation merely recites an obvious aspect of digital logic design. As explained above at IX.B.1, Bruckman's equipment 22 is a two-level structure that resembles a two-level demultiplexer. Bruckman's concentrators 50 act as first-level demultiplexers as each concentrator 50 selects one of four links 30

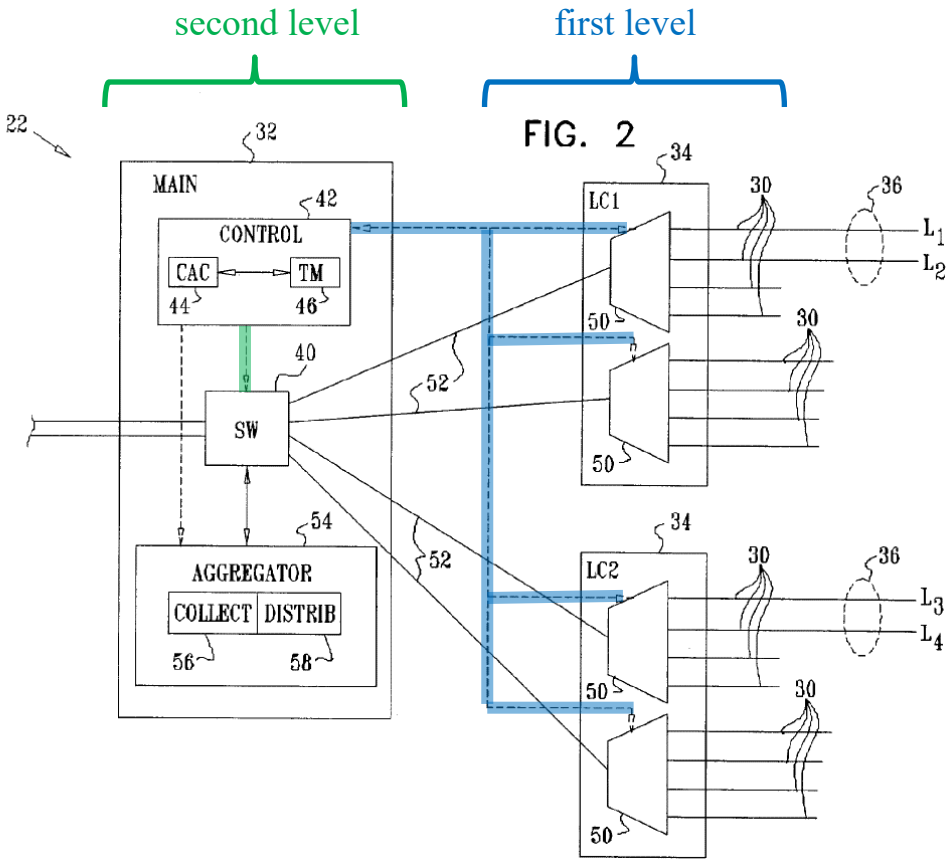
to transmit data received from a single trace 52. Bruckman's switching core 40 acts as the second-level demultiplexer by distributing data frames from the input channel amongst four separate traces 52.



Bruckman, Fig. 2 (annotated); Ex.1003, ¶137.

Given this well-known, two-level structure, a POSITA would have found it obvious for Bruckman's Equipment 22 to be controlled in a manner consistent with well-known two-level demultiplexer control techniques. Ex.1003, ¶138. Indeed, Bruckman's equipment 22 includes a control line from the controller 42 to the

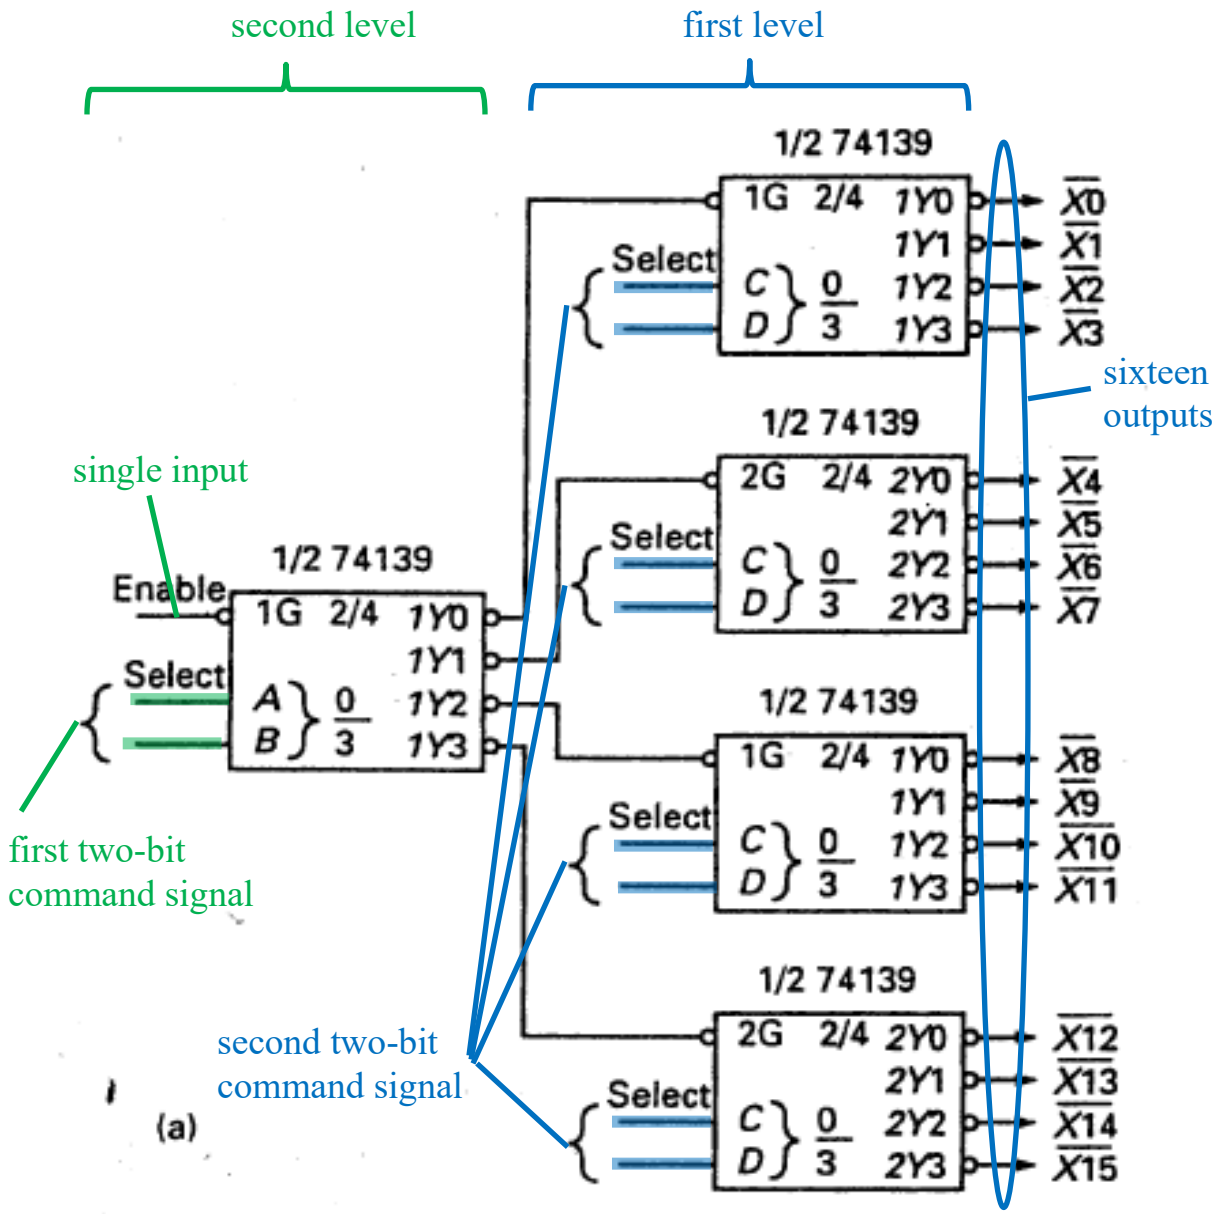
switching core (second-level demultiplexer) shown below in green. Equipment 22 also includes control lines from the controller 42 to each of the concentrators (first-level demultiplexers) shown below in blue.



Bruckman, Fig. 2 (annotated); Ex.1003, ¶138.

Bruckman's description of these control lines is limited. Because each of the demultiplexing devices (switching core 40 and concentrators 50) in Fig.2 has one input and four outputs, a POSITA would have found it obvious to control each with a two-bit command signal. Ex.1003, ¶139. A two-bit command signal would efficiently be able to specify one of four operating states, each corresponding to the

four possible outputs. Ex.1003, ¶139. Thus, in Bruckman's example, a two-bit command signal would be used to control the switching core (second-level demultiplexer) to select a trace 52. A different two-bit signal would be used to control the concentrators (first-level demultiplexers) to select a link 30, consistent with known techniques for controlling two-level demultiplexers. Such techniques, illustrated below, would have been in the background knowledge of a POSITA. *See supra* IX.B.5; Ex.1003, ¶139.



Ex.1007, Fig. 5.17(a) (annotated); Ex.1003, ¶139.

As explained above at [1.6], Bruckman's hashing function produces a single hash value that is used to select both a line from the first level and a line from the second level. In Bruckman's example where there are 16 possible links, the lagSize parameter is 16. The operation "hash % 16" would produce a value ("*modulo*")

within the range of 0 and 15. Ex.1003, ¶140. The sixteen different values between 0 and 15 are represented in binary as a four-bit value ("*binary representation of the modulo*").

Consistent with a POSITA's background knowledge of using two two-bit control signals to control a two-level demultiplexer (formed from 1:4 demultiplexers), two of the four hash bits would be used to control the first level and the remaining two of the four hash bits would be used to control the second level. *See* Ex.1007, Fig. 5.17(a); Ex.1003, ¶141. Thus, a POSITA would have found it obvious for Bruckman's control signals (represented by dotted lines in Fig. 2) to be two-bit signals formed from different bits of the hash value. Ex.1003, ¶141.

Thus, because Bruckman illustrates control signals to both levels of a two-level demultiplexer, and a POSITA would have known the technique for controlling such a structure involves subsets of bits for control of different levels, Bruckman renders obvious "*selecting the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo*" as claimed. Ex.1003, ¶¶137-42.

17. Claim 12

[12.1] ***The method according to claim 1, wherein the at least one of the frame attributes comprises at least one of a layer 2 header field, a layer 3 header field, a layer 4 header field, a source Internet Protocol (IP) address, a destination IP***

address, a source medium access control (MAC) address, a destination MAC address, a source Transmission Control Protocol (TCP) port and a destination TCP port.

As explained above at [1.5], Bruckman's equipment 22 receives a data frame with frame attributes. As noted in Bruckman's Background section, it was known to use "information carried in each Ethernet frame in order to make [a] decision as to the physical port to which the frame should be sent.… The information used to assign conversations to ports could thus include one or more of the following pieces of information: **a) Source MAC address b) Destination MAC address**." Ex.1005, [0005]-[0011]. Moreover, POSITAs would have recognized that the list of options recited in claim 12 are well-known data fields and that are commonly part of a layer 2 data frame. *See e.g.*, Ex.1010; Ex.1003, ¶143.

Thus, because Bruckman's equipment receives a data frame having at least a source MAC address and a destination MAC address, Bruckman renders this limitation obvious. Ex.1003, ¶¶143-45.

18. Claim 13

[13.0]-[13.5] ***A method for … and second physical links,***

See [1.0]-[1.7]. Ex.1003, ¶¶146-51.

[13.6] ***coupling the network node to the one or more interface modules and coupling each of the one or more interface modules to the communication network comprising specifying bandwidth requirements comprising at least one of a committed information rate (CIR), a peak information rate (PIR) and an excess information rate (EIR) of a communication service provided by the***

communication network to the network node, and

First, as explained above at [1.1] and [1.3], Bruckman renders obvious "*coupling the network node to the one or more interface modules and coupling each of the one or more interface modules to the communication network*."

Second, Bruckman explains that service level agreements between network service providers (including a "*communication network*") and customers (including a "*customer node*") specify bandwidth in various terms. "In general, the bandwidth guaranteed by a service provider, referred to as the peak information rate (PIR), may include either CIR, or EIR, or both CIR and EIR (in which case PIR=CIR+EIR)." Ex.1005, [0013].

Thus, because Bruckman describes that providing guaranteed bandwidth in terms of either CIR, or EIR, or both CIR and EIR (in which case PIR=CIR+EIR), Bruckman renders this limitation obvious. Ex.1003, ¶¶152-54.

[13.7] ***allocating a bandwidth for the communication service over the first and second physical links responsively to the bandwidth requirements.***

Bruckman describes allocating bandwidth in a manner so as to exceed the guaranteed bandwidth ("*responsively to the bandwidth requirements*"). *See* Ex.1005, [0030]. Bruckman describes "a method for establishing a connection with a guaranteed bandwidth for transmitting data between first and second endpoints," Ex.1005, [0027]; *see also* [0016], [0030].

Accordingly, establishing the communication sessions in Bruckman are for the purpose of providing the guaranteed bandwidth and is thus "*responsively to the bandwidth requirements*."

Thus, because Bruckman allocates bandwidth to exceed the guaranteed bandwidth, Bruckman renders this limitation obvious. Ex.1003, ¶¶155-57.

19. Claim 14

[14.0] ***A method for connecting user ports to a communication network, comprising:***

As explained above at [1.0], [1.1], and [1.3], Bruckman describes equipment 22 that connects a plurality of links 30 (which may correspond to the claimed "*ports*") to a communication network. Alternatively, Bruckman's links are associated with ports. Ex.1005, [0056] ("multiple ports that serve respective links 30"). The ports associated with links 30 correspond to "*user ports*" because they connect to user nodes, as explained in [2.2]. Thus, Bruckman renders this limitation obvious. Ex.1003, ¶158.

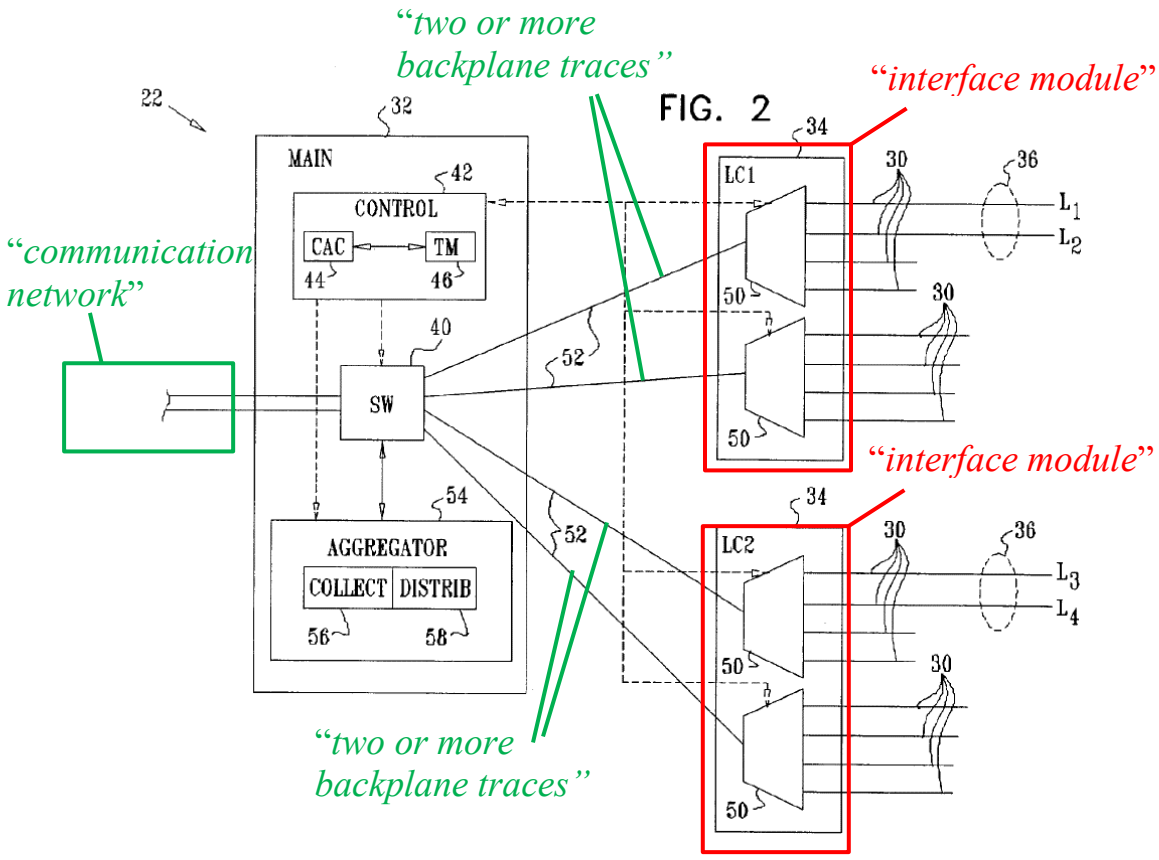
[14.1] ***coupling the user ports to one or more user interface modules;***

Consistent with the discussion at [1.1] and [14.0], Bruckman renders obvious "*coupling the user ports to one or more user interface modules*" as claimed. Ex.1003, ¶159.

[14.2] ***coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel,***

First, consistent with the discussion at [1.3], Bruckman renders obvious "coupling each user interface module to the communication network via a backplane" as claimed. *See also* [3.2] (explaining that Bruckman's traces 52 are backplane traces).

Second, as shown in Fig. 2 below, each of Bruckman's two line cards 34 ("*interface modules*") includes at least two traces 52 that connect that line card to the communication network.



Bruckman, Fig. 2 (annotated); Ex.1003, ¶161.

Thus, Bruckman renders obvious "*coupling each user interface module to*

the communication network via a backplane using two or more backplane traces arranged in parallel" as claimed. Ex.1003, ¶¶160-62.

[14.3] ***at least one of said backplane traces being bi-directional and operative to communicate in both an upstream direction and a downstream direction;***

Consistent with the discussion above at [1.4], Bruckman renders this limitation obvious. Ex.1003, ¶163.

[14.4] ***receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;***

Consistent with the discussion above at [1.6], Bruckman renders this limitation obvious. Ex.1003, ¶164.

[14.5] ***for each data frame, selecting responsively to at least one of the respective frame attributes a backplane trace from the two or more backplane traces; and***

Consistent with the discussion above at [1.6], Bruckman renders this limitation obvious. *See also* [3.1] (explaining that Bruckman's traces 52 are backplane traces). Ex.1003, ¶165.

[14.6] ***sending the data frame over the selected backplane trace; said sending comprising communicating along said at least one of said backplane traces.***

Consistent with the discussion above at [1.7], Bruckman renders this limitation obvious. *See also* [3.1] (explaining that Bruckman's traces 52 are backplane traces). Ex.1003, ¶166.

20. Claim 15

[15.0]-[15.4] ***A method for … backplane traces; and***

See [14.0]-[14.5]. Ex.1003, ¶¶167-71.

[15.5] ***sending the data frame over the selected backplane trace, at least some of the backplane traces being aggregated into an Ethernet link aggregation (LAG) group.***

First, consistent with the discussion above at [1.7], Bruckman renders obvious "*sending the data frame over the selected backplane trace*" as claimed. *See also* [3.1] (explaining that Bruckman's traces 52 are backplane traces).

Second, Bruckman explains that the traces 52 are part of a link aggregation group: "As a result of spreading group 36 over two (or more) line cards, **the link aggregation function applies** not only to links 30 in group 36 but also **to traces 52 that connect to multiplexers 50** that serve these links." Ex.1005, [0057]. As also explained above at [4.6], Bruckman's link aggregation groups may be Ethernet link aggregation groups.

Thus, because Bruckman's equipment 22 sends data over selected traces 52 and links 30, and the traces 52 are part of an Ethernet link aggregation group, Bruckman renders this limitation obvious. Ex.1003, ¶¶172-74.

21. Claim 16

[16.1] ***The method according to claim 14, wherein selecting the backplane trace comprises applying a hashing function to the at least one of the frame attributes.***

See claims 8 and 9. Ex.1003, ¶175.

22. Claim 17

[17.0] ***Apparatus for connecting a network node with a communication network, comprising:***

As explained above at [1.1] and [1.3], Bruckman describes Equipment 22 ("*apparatus*") that connects customer nodes (a "*network node*") with a communication network. Thus, Bruckman renders this limitation obvious. Ex.1003, ¶176.

[17.1] ***one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network,***

First, Bruckman's Equipment 22 includes a plurality of line cards 34 ("*one or more interface modules*"). *See* Ex.1005, [0049].

Second, Bruckman's line cards process data frames between a customer node ("*network node*") and the communication network. "Each line card 34 comprises one or more concentrators 50, which comprise multiple ports that serve respective links 30. The concentrators multiplex data traffic between links 30 and traces 52, which connect the concentrators to switching core 40." Ex.1005, [0056].

Third, consistent with the discussion at [1.5], Bruckman's data frames include frame attributes.

Thus, because Bruckman's equipment includes a plurality of line cards that multiplex traffic from a customer node to a communication network, and the traffic

includes data frames with frame attributes, Bruckman renders obvious "*one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network*" as claimed. Ex.1003, ¶¶177-80.

[17.2] ***at least one of said interface modules being operative to communicate in both an upstream direction and a downstream direction;***

Consistent with the discussion above at [1.2] and [1.4], Bruckman's Equipment processes traffic in both upstream and downstream directions. Accordingly, the line cards ("*interface modules*") are "*operative to communicate in both an upstream direction and a downstream direction.*" Thus, Bruckman renders this limitation obvious. Ex.1003, ¶181.

[17.3] ***a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules*;**

Consistent with the discussion above at [1.1], Bruckman renders this limitation obvious. Ex.1003, ¶182.

[17.4] ***a second group of second physical links arranged in parallel so as to couple the one or more interface modules to the communication network; and***

Consistent with the discussion above at [1.3], Bruckman renders this limitation obvious. Ex.1003, ¶183.

[17.5] ***a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which***

to send the data frame;

First, Bruckman's Equipment 22 includes a controller 42 ("*control module*"): "The operation of switch 40 is managed by a controller 42, typically an embedded microprocessor with suitable software for carrying out the functions described herein." Ex.1005, [0049]. Bruckman further describes additional software components running on controller 42, including "Connection Admission Control entity (CAC) 44," "traffic manager 46," and "Aggregator 54." Ex.1005, [0050], [0056]-[0057].

Second, consistent with the discussion above at [1.6], Bruckman's controller 42 (including aggregator 54) is "*arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame*" as claimed. Thus, Bruckman (alone or in combination with Basso) renders this limitation obvious. Ex.1003, ¶¶184-85.

[17.6] *…being bi-directional links….*

See [1.2] and [1.4]. Ex.1003, ¶186.

23. Claim 18

[18.1] *The apparatus according to claim 17, and comprising a backplane to which the one or more interface modules are coupled, wherein the second physical links comprise backplane traces formed on the backplane.*

Consistent with the discussion above at [3.1], Bruckman renders this limitation obvious. Ex.1003, ¶187.

24. Claim 19

[19.0]-[19.4] ***Apparatus for … send the data frame,***

See [17.0]-[17.5]. Ex.1003, ¶¶188-92.

[19.5] ***at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.***

See [4.6]. Ex.1003, ¶193.

25. Claim 20

[20.0]-[20.4] ***Apparatus for … send the data frame,***

See [17.0]-[17.5]. Ex.1003, ¶¶194-98.

[20.5] ***…links being aggregated into an external Ethernet link aggregation (LAG) group….***

See [5.6] and [5.7]. Ex.1003, ¶199.

26. Claims 21-27

Claims 21-27 depend (directly or indirectly) from apparatus claim 17 and recite in apparatus form substantially the same subject matter limitations as claims 6-12 (respectively), which depend from method claim 1. Accordingly, claims 21-27 are obvious for the same reasons discussed previously for claims 6-12, respectively. Ex.1003, ¶¶200-10.

27. Claim 28

[28.0]-[28.4] ***Apparatus for … send the data frame,***

See [17.0]-[17.5]. Ex.1003, ¶¶211-15.

[28.5] ***the communication network being arranged to provide a communication service…***

See [13.6]. Ex.1003, ¶216

[28.6] ***the first and second groups of physical links being dimensioned to provide an allocated bandwidth for the communication service responsively to the bandwidth requirements.***

Consistent with the discussion above at [13.7], Bruckman renders this limitation obvious. A POSITA would have understood that if the links can provide the allocated bandwidth, they are thus "*dimensioned*" to provide such bandwidth. The links would otherwise not be able to provide the required bandwidth. Ex.1003, ¶217.

28. Claim 29

[29.0] ***Apparatus for ... modules being bi-directional…;***

See [17.0]-[17.2]. Ex.1003, ¶¶218-20.

[29.3] ***a backplane having the one or more user interface modules coupled thereto and comprising a plurality of backplane traces arranged in parallel so as to transfer the data frames between the one or more user interface modules and the communication network,***

Consistent with the discussion above at [3.1] and [18.1], Bruckman renders this limitation obvious. Ex.1003, ¶221.

[29.4] ***at least one of said backplane traces being bi-directional…; and***

Consistent with the discussion above at [1.2] and [1.4], Bruckman renders this limitation obvious. Ex.1003, ¶222.

[29.5] ***a control module,….***

See [19.4]. Ex.1003, ¶223.

29. Claim 30

[30.0] ***Apparatus for ... comprising:***

See [17.0] and [14.0]. Ex.1003, ¶224.

[30.1] ***one or more …;***

Consistent with the discussion above at [17.1], Bruckman renders this limitation obvious. Ex.1003, ¶225.

[30.2] ***a backplane having the one or more user interface modules…;***

See [29.3]. Ex.1003, ¶226.

[30.4] ***a control module,…;***

See [19.4]. Ex.1003, ¶227.

[30.5] ***at least some of the backplane traces are aggregated into an Ethernet link aggregation (LAG) group***.

See [15.5]. Ex.1003, ¶228.

30. Claim 31

[31.1] ***The apparatus ... backplane trace.***

See claims 8 and 9. Ex.1003, ¶229.

C. Grounds 3 & 4: Claims 11 and 26 are obvious under 35 U.S.C. § 103(a) over Bruckman (alone or with Basso) in view of Holdsworth.

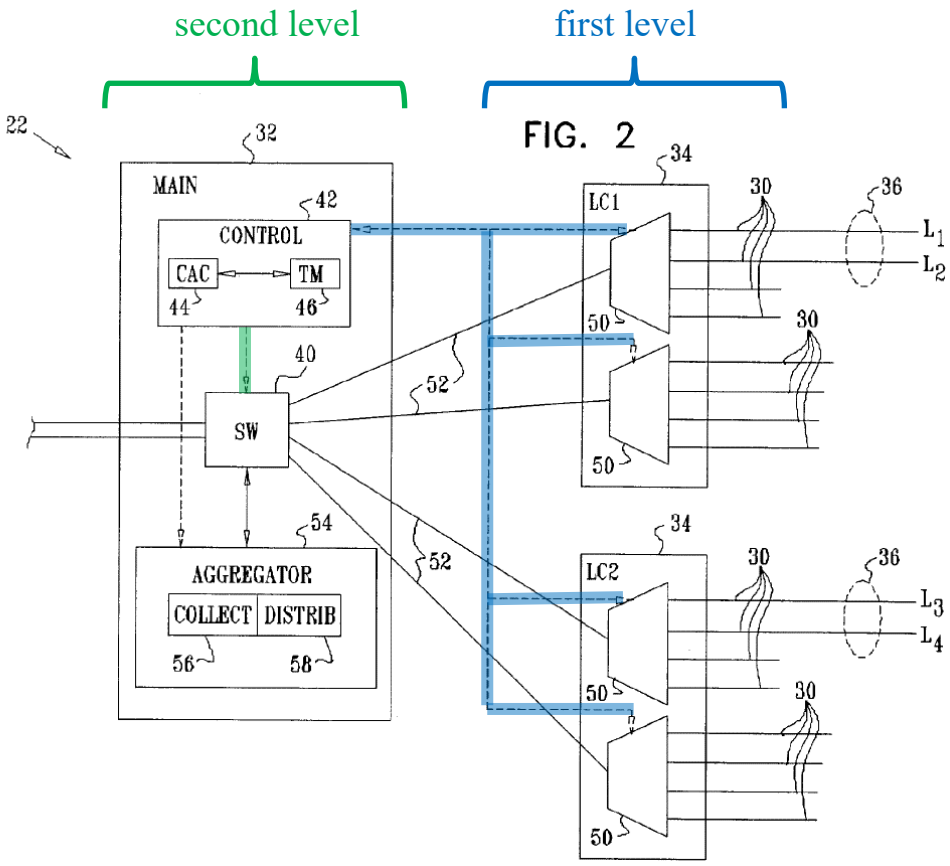
Grounds 1 and 2 show how Bruckman renders obvious limitations related to using different subsets of bits to select different levels in a multi-level demultiplexer. Grounds 1 and 2 further rely on the Holdsworth textbook as evidence of the background knowledge of a POSITA. To the extent Patent Owner argues that such digital logic design would not have been within the background knowledge of a POSITA, Grounds 3 and 4 are presented to show that combining the disclosures of Holdsworth with the other prior art would have been obvious to a POSITA. Ex.1003, ¶230.

1. Reasons to Combine Holdsworth with Bruckman

A POSITA would have found it obvious to use basic digital logic design to implement the control signals for Bruckman's two-level demultiplexer. For example, a POSITA would have found it obvious to use a first subset of bits for a first-level demultiplexer, and a second subset of bits for the second-level demultiplexers, as evidenced by Holdsworth. Ex.1003, ¶231.

As explained above, Bruckman's Equipment 22 uses a two-level demultiplexer structure for distributing incoming packets from the communication network among the links 30. Bruckman's depiction of two-level demultiplexer includes dotted lines from the controller to the first-level and second-level

demultiplexers, as shown in Fig. 2 below.



Ex.1005, Fig. 2 (annotated); Ex.1003, ¶232.

A POSITA would have recognized that these dotted lines represent command signals for controlling each of the demultiplexers (switching core 40 and concentrators 50). Ex.1003, ¶233. Bruckman illustrates these control lines but provides few implementation details about the command signals used on the control lines, indicating that such implementation would have been within the knowledge and skill of a POSITA. Ex.1003, ¶233. Accordingly, a POSITA looking at the teachings of Bruckman would have looked to known implementation

techniques for the command signals of Bruckman's two-level demultiplexer. Ex.1003, ¶233. Bruckman explains that "at least some of the functions of the aggregator may be carried out by hard-wired logic or by a programmable logic component, such as a gate array." Ex.1005, [0057]. Accordingly, a POSITA would have looked to known logic design techniques. Ex.1003, ¶233.

Holdsworth represents basic logic design that is applicable to Bruckman's two-level demultiplexer structure. As stated in the preface, Holdsworth is intended to "act as a reference text for graduates working in this field." Holdsworth, preface. A POSITA would have had a reasonable expectation of success when using Holdsworth's technique because Holdsworth provides an example that is similar to that of Bruckman's example—a plurality of 1:4 demultiplexers that form a 1:16 demultiplexer. A POSITA would have therefore expected that using a two-bit command signal to control a first-level demultiplexer, and a second two-bit command signal to control second-level demultiplexers would have worked as well in Bruckman's two-level demultiplexer. Ex.1003, ¶234.

Given that Bruckman's hash function would produce a hash value in the range of zero to fifteen (and thus be readily represented as a four-bit value), a POSITA would have found it obvious to use two of the four bits to control the second-level demultiplexer (i.e., the switching core) and the other two of the four bits to control the first-level demultiplexers (i.e., the concentrators). Ex.1003, ¶235.

Regarding Ground 4, application of Holdsworth technique is also consistent with Basso's technique of using the single hash computation to select both a second-level line (one of the traces 52) and a first-level line (one of the links 30). This would have yielded the predictable result of effectively and efficiently controlling both levels of the two-level demultiplexer using the four-bit hash computation described by Bruckman.

Thus, the combination of Holdsworth with Bruckman represents application of a known technique (Holdsworth's different bit subsets for different levels of a two-level demultiplexer) to a known method (Bruckman's two-level demultiplexer) to yield predictable results (controlling both levels with a single hash computation). Ex.1003, ¶¶231-36.

2. Claims 11 and 26

The claim analysis for this ground is substantially the same as the analysis provided in Grounds 1 and 2. *See supra* IX.B.16, IX.B.26; Ex.1003, ¶237.

X. DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE

A. Discretionary denial under the *Fintiv* factors is not appropriate

The six factors considered for § 314 denial strongly favor institution. *See Apple Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 (PTAB Mar. 20, 2020) (precedential).

1. No evidence regarding a stay

No motion to stay has been filed, so the Board should not infer the outcome of such a motion. *Sand Revolution II LLC v. Continental Intermodal Group – Trucking LLC*, IPR2019-01393, Paper 24 at 7 (PTAB 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 on discretionary denial.

2. Parallel proceeding trial date

This factor weighs strongly against discretionary denial because the projected trial date—based on median time-to-trial statistics—is in August 2024, after the Board's Final Written Decision is expected in July 2024.[4] While trial is currently proposed for March 4, 2024 (Ex.1014), the Board recognizes "that scheduled trial dates are unreliable and often change." *See* Director's June 21, 2022 Memorandum on Discretionary Denials ("Memo"), 8. Accordingly, the Board now uses median time-to-trial statistics in the relevant venue to determine a projected trial date for Fintiv purposes. Memo, 9.

[4] July 2024 is 18 months after January 2023, when Petitioner expects a notice of accorded filing date for this petition.

The co-pending district court case was filed in the Eastern District of Texas on July 22, 2022. *See* Ex.1012. The current median time-to-trial in the Eastern District of Texas is 24.5 months. Ex.1013, 5. Accordingly, the projected trial date for *Fintiv* purposes is August of 2024—approximately 24 months after July 2022, and after the Board's Final Written Decision is expected in July of 2024. Because the projected trial date is "around the same time or after" the Board's expected final written decision, this factor weighs in favor of institution.

3. Investment in the parallel proceeding

The co-pending litigation is in its very early stages, and the investment in it has been minimal. The parties have not exchanged preliminary positions on claim construction or invalidity, expert discovery has not begun, and the parties have not exchanged their first set of discovery requests. *See PEAG LLC v. Varta Microbattery GmbH*, IPR2020-01214, Paper 8, 17 (Jan. 6, 2021). Further, the Markman hearing is not scheduled until September of 2023, two months after an expected institution decision by the Board. Ex.1014, 3.

Moreover, Petitioner only learned which claims were being asserted on November 3, 2022. *See* Ex.1015. Under *Fintiv*, Petitioner's prompt filing "weigh[s] against exercising the authority to deny institution." *Fintiv*, Paper 11 at 11 ("If the evidence shows that the petitioner filed the petition expeditiously, such as promptly after becoming aware of the claims being asserted, this fact has

weighed against exercising the authority to deny institution under NHK"). This factor favors institution.

4. Overlapping issues with the parallel proceeding

There is no present overlap of prior art issues due to the early stage of district court litigation. For example, Petitioner has not served its preliminary invalidity contentions in the district court proceeding. Consequently, this factor favors institution.

5. Identity of parties

Petitioner is a defendant in the litigation. That is true of most Petitioners in IPR proceedings. Accordingly, this factor should not be a basis for denying institution.

6. Other circumstances

The prior art presented in this Petition renders the Challenged Claims unpatentable as obvious. "[T]he PTAB will not deny institution of an IPR or PGR under Fintiv (i) when a petition presents compelling evidence of unpatentability." Memo, 2. "Compelling, meritorious challenges are those in which the evidence, if unrebutted in trial, would plainly lead to a conclusion that one or more claims are unpatentable by a preponderance of the evidence." Memo, 4. Here, the petition plainly shows that the '740 patent claims no more than known concepts of bi-

directional link aggregation. The evidence of unpatentability is compelling, and thus the PTAB should not deny institution under Fintiv.

As such, because the *Fintiv* factors are either neutral or weigh against discretionary denial, and institution should not be denied on discretionary factors.

B. Discretionary denial under 35 U.S.C. § 325(d) is not appropriate

Denial under § 325(d) is not warranted because the challenges presented in this petition are neither cumulative nor redundant to the prosecution of the '740 Patent. To begin with, this petition relies in part on Basso and Holdsworth, neither of which was considered in prosecution. And while Bruckman was disclosed to the examiner after most claims had already been allowed, the Examiner erred by overlooking Bruckman's highly relevant teachings. *See* Ex.1002, 44. This was "material" error because the overlooked teachings of Bruckman render obvious the limitations that the Applicant had added to gain allowance. Discretionary denial is therefore not appropriate. *Advanced Bionics, LLC v. MED-EL Elektromedizinische Geräte GmbH*, IPR2019-01469, Paper 6 (PTAB Feb. 13, 2020) (precedential) ("*Advanced Bionics*").

Under the first prong of the *Advanced Bionics* framework, Petitioner acknowledges that Bruckman was "previously presented to the Office." *Advanced Bionics* at 7-8. Under the second prong, however, the evidence shows that the Examiner "erred in a manner material to the patentability of challenged claims."

Id. The second prong is guided by *Becton, Dickinson* factors (c), (e), and (f):

- (c) the extent to which the asserted art was evaluated during examination, including whether the prior art was the basis for rejection;
- (e) whether petitioner has pointed out sufficiently how the examiner erred in its evaluation of the asserted prior art; and
- (f) the extent to which additional evidence and facts presented in the petition warrant reconsideration of the prior art or arguments.

Id. at 9-10, n. 10 (citing *Becton, Dickinson & Co. v. B. Braun Melsungen AG*, IPR2017-01586, Paper 8 at 17-18 (Dec. 15, 2017)). These factors weigh against exercising discretion.

1. *Becton, Dickinson* Factor (c)

Bruckman was never cited in a rejection. Bruckman was introduced to the record after the Examiner had allowed several dependent claims, and after the Applicant had amended various independent claims to gain allowance. Just days after the examiner signed an IDS listing Bruckman, the Office mailed a Notice of Allowance for all claims with not statement of reasons for allowance. Ex.1002, 40, 44. Thus, the examiner made a "material error" by overlooking the disclosure of Bruckman that teaches the very concepts related to the bi-directional nature of the links that were added to the claims to gain allowance. Bruckman also teaches the very concepts that were in the dependent claims that examiner initially allowed. The sparse references to Bruckman in the file history indicate that Bruckman

received—at most—only insufficient, cursory attention during prosecution. A reference merely being of record is not sufficient reason to exercise discretion. *See Navistar, Inc. v. Fatigue Fracture Tech., LLC*, IPR2018-00853, Paper 13 at 17 (Sept. 12, 2018) ("Under [Becton factors] (c), (d), and (f) . . . the fact that [references] were of record, but not applied in any rejection by the Examiner . . . provides little impetus for us to exercise our discretion to deny institution under § 325(d).")

Accordingly, because the Office overlooked the teachings of Bruckman when allowing the '740 patent, and because Basso and Holdsworth have never been previously presented by the Office, factor (c) favors institution.

2. *Becton, Dickinson* Factors (e) and (f):

Advanced Bionics explains that "if the record of the Office's previous consideration of the art is not well developed or silent, then a petitioner may show the Office erred by overlooking something persuasive under factors (e) and (f)." *Advanced Bionics* at 10. Here, the Office's underdeveloped consideration of Bruckman was "material error" because it overlooked Bruckman's description of how data is both passed from the communication network to the customer nodes (*See* Ex.1005, [0058]) and passed from the customer nodes to the communication network (Ex.1005, [0065]). Bruckman also explains that the links may be "full-duplex" Ethernet links, which a POSITA would have understood to mean that they

are capable of transmitting data in both directions simultaneously. Ex.1005, [0047]; Ex.1009. By allowing claims based on limitations that were taught in the cited art of record—including the specific limitations the examiner had previously deemed allowable—the Office overlooked teachings that had a significant impact on patentability. *Advanced Bionics* at 8 n.9 ("material error may include misapprehending or overlooking specific teachings of the relevant prior art where those teachings impact patentability of the challenged claims"). This petition's grounds of unpatentability are not merely "a disagreement with a specific finding of record by the Office." *Advanced Bionics* at 10-11. Rather, this petition relies upon teachings in Bruckman that were simply overlooked during prosecution.

The Board has consistently declined to discretionarily deny institution under similar facts. *See, e.g., DISH Network LLC v. Sound View Innovations, LLC*, IPR2020-01041, Paper 13, 22 (PTAB Jan 19, 2021) (granting institution where "Petitioner has sufficiently shown Examiner error by pointing out specific teachings of the [previously-asserted] prior art that the Examiner overlooked"); *Trans Ova Genetics, LC v. XY, LLC*, IPR2018-00250, Paper 9, 18-19 (PTAB June 27, 2018) (granting institution where "the Examiner manifestly failed to appreciate [disclosures]…from [the previously-cited art]").

With respect to factor (f), the Petition is supported by an expert declaration (Ex.1003) by Dr. Houh explaining how a POSITA would have understood

Bruckman and the obviousness of combining its teachings with those of Basso and Holdsworth. *See generally*, Ex.1003; *see also* Ex.1004. This new evidence also weighs against discretionary denial. *See, e.g., Puma N. Am., Inc. v. Nike, Inc.*, IPR2019-01058, Paper 10, 19 (PTAB Oct. 31, 2019) (instituting where petition presented "new non-cumulative evidence…. probative to issues of patentability and helpful to our consideration of a prior art combination that was not before the Examiner").

Accordingly, the Board should not exercise its discretion to deny this petition under § 325(d).

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

The '740 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 (PTAB Sept. 6, 2016) (Section II.B.4.i. precedential).

XI. CONCLUSION

Accordingly, Petitioner has established a reasonable likelihood that the Challenged Claims are unpatentable.

Respectfully submitted,

Dated: January 9, 2023

/Theodore M. Foster/

HAYNES AND BOONE, LLP
2323 Victory Avenue, Suite 700
Dallas, Texas 75219
Customer No. 27683

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

XII. 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 '740 patent is involved in the following case:

Case Heading	Number	Court	Date
Orckit Corporation v. Cisco Systems, Inc.	2:22-cv-276-JRG-RSP	EDTX	July 22, 2022

C. Lead and Back-up Counsel and Service Information

Lead Counsel
Theodore M. Foster
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219
Phone: (303) 382-6205
Fax: (214) 200-0853
ipr.theo.foster@haynesboone.com
USPTO Reg. No. 57,456

Back-up Counsel
David L. McCombs
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219
Phone: (214) 651-5533
Fax: (214) 200-0853
david.mccombs.ipr@haynesboone.com
USPTO Reg. No. 32,271

Gregory P. Huh
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219
Phone: (972) 739-6939
Fax: (214) 200-0853
gregory.huh.ipr@haynesboone.com
USPTO Reg. No. 70,480

Calmann J. Clements
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219

Phone: (972) 739-8638
Fax: (214) 200-0853
calmann.clements.ipr@haynesboone.com
USPTO Reg. No. 66,910

Please address all correspondence to lead and back-up counsel. Petitioner consents to service in this proceeding by email at the addresses above.

XIII. CLAIMS APPENDIX

[1.0] 1. A method for communication, comprising:

[1.1] coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel,

[1.2] at least one of said first physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction;

[1.3] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel,

[1.4] at least one of said second physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction;

[1.5] receiving a data frame having frame attributes sent between the communication network and the network node;

[1.6] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and

[1.7] sending the data frame over the selected first and second physical links, said sending comprising communicating along at least one of said bi-directional links.

[2.1] 2. The method according to claim 1, wherein the network node comprises a user node, and

[2.2] wherein sending the data frame comprises establishing a communication service between the user node and the communication network.

[3.1] 3. The method according to claim 1, wherein the second physical links comprise backplane traces formed on a backplane to which the one or more interface modules are coupled.

[4.0] 4. A method for communication, comprising:

[4.1] coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel;

[4.2] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel;

[4.3] receiving a data frame having frame attributes sent between the communication network and the network node;

[4.4] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and

[4.5] sending the data frame over the selected first and second physical links,

[4.6] at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.

[5.0] 5. A method for communication, comprising:

[5.1] coupling a network node to one or more interface modules using a first

group of first physical links arranged in parallel;

[5.2] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel;

[5.3] receiving a data frame having frame attributes sent between the communication network and the network node;

[5.4] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and

[5.5] sending the data frame over the selected first and second physical links,

[5.6] coupling the network node to the one or more interface modules comprises aggregating two or more of the first physical links into an external Ethernet link aggregation (LAG) group

[5.7] so as to increase a data bandwidth provided to the network node.

[6.1] 6. The method according to claim 1, wherein coupling each of the one or more interface modules to the communication network comprises at least one of multiplexing upstream data frames sent from the network node to the communication network, and

[6.2] demultiplexing downstream data frames sent from the communication network to the network node.

[7.1] 7. The method according to claim 1, wherein selecting the first and second

physical links comprises balancing a frame data rate among at least some of the first and second physical links.

[8.1] 8. The method according to claim 1, wherein selecting the first and second physical links comprises applying a mapping function to the at least one of the frame attributes.

[9.1] 9. The method according to claim 8, wherein applying the mapping function comprises applying a hashing function.

[10.1] 10. The method according to claim 9, wherein applying the hashing function comprises determining a hashing size responsively to a number of at least some of the first and second physical links,

[10.2] applying the hashing function to the at least one of the frame attributes to produce a hashing key,

[10.3] calculating a modulo of a division operation of the hashing key by the hashing size, and

[10.4] selecting the first and second physical links responsively to the modulo.

[11.1] 11. The method according to claim 10, wherein selecting the first and second physical links responsively to the modulo comprises selecting the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo.

[12.1] 12. The method according to claim 1, wherein the at least one of the frame

attributes comprises at least one of a layer 2 header field, a layer 3 header field, a layer 4 header field, a source Internet Protocol (IP) address, a destination IP address, a source medium access control (MAC) address, a destination MAC address, a source Transmission Control Protocol (TCP) port and a destination TCP port.

[13.0] 13. A method for communication, comprising:

[13.1] coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel;

[13.2] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel;

[13.3] receiving a data frame having frame attributes sent between the communication network and the network node;

[13.4] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and

[13.5] sending the data frame over the selected first and second physical links,

[13.6] coupling the network node to the one or more interface modules and coupling each of the one or more interface modules to the communication network comprising specifying bandwidth requirements comprising at least one of a committed information rate (CIR), a peak information rate (PIR)

and an excess information rate (EIR) of a communication service provided by the communication network to the network node, and

[13.7] allocating a bandwidth for the communication service over the first and second physical links responsively to the bandwidth requirements.

[14.0] 14. A method for connecting user ports to a communication network, comprising:

[14.1] coupling the user ports to one or more user interface modules;

[14.2] coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel,

[14.3] at least one of said backplane traces being bi-directional and operative to communicate in both an upstream direction and a downstream direction;

[14.4] receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;

[14.5] for each data frame, selecting responsively to at least one of the respective frame attributes a backplane trace from the two or more backplane traces; and

[14.6] sending the data frame over the selected backplane trace; said sending comprising communicating along said at least one of said backplane traces.

[15.0] 15. A method for connecting user ports to a communication network, comprising:

[15.1] coupling the user ports to one or more user interface modules;

[15.2] coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel;

[15.3] receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;

[15.4] for each data frame, selecting responsively to at least one of the respective frame attributes a backplane trace from the two or more backplane traces; and

[15.5] sending the data frame over the selected backplane trace, at least some of the backplane traces being aggregated into an Ethernet link aggregation (LAG) group.

[16.1] 16. The method according to claim 14, wherein selecting the backplane trace comprises applying a hashing function to the at least one of the frame attributes.

[17.0] 17. Apparatus for connecting a network node with a communication network, comprising:

[17.1] one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network,

[17.2] at least one of said interface modules being operative to communicate in

both an upstream direction and a downstream direction;

[17.3] a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;

[17.4] a second group of second physical links arranged in parallel so as to couple the one or more interface modules to the communication network; and

[17.5] a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame;

[17.6] at least one of said first physical links and at least one of said second links being bi-directional links operative to communicate in both said upstream direction and said downstream direction.

[18.1] 18. The apparatus according to claim 17, and comprising a backplane to which the one or more interface modules are coupled, wherein the second physical links comprise backplane traces formed on the backplane.

[19.0] 19. Apparatus for connecting a network node with a communication network, comprising:

[19.1] one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the

communication network;

[19.2] a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;

[19.3] a second group of second physical links arranged in parallel so as to couple the one or more interface modules to the communication network; and

[19.4] a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame,

[19.5] at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.

[20.0] 20. Apparatus for connecting a network node with a communication network, comprising:

[20.1] one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network;

[20.2] a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;

[20.3] a second group of second physical links arranged in parallel so as to couple

the one or more interface modules to the communication network; and

[20.4] a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame,

[20.5] two or more of the first physical links being aggregated into an external Ethernet link aggregation (LAG) group so as to increase a data bandwidth provided to the network node.

[21.1] 21. The apparatus according to claim 17, and comprising a multiplexer, which is arranged to perform at least one of multiplexing upstream data frames sent from the network node to the communication network, and

[21.2] demultiplexing downstream data frames sent from the communication network to the network node.

[22.1] 22. The apparatus according to claim 17, wherein the control module is arranged to balance a frame data rate among at least some of the first and second physical links.

[23.1] 23. The apparatus according to claim 17, wherein the control module is arranged to apply a mapping function to the at least one of the frame attributes so as to select the first and second physical links.

[24.1] 24. The apparatus according to claim 23, wherein the mapping function comprises a hashing function.

[25.1] 25. The apparatus according to claim 24, wherein the control module is arranged to determine a hashing size responsively to a number of at least some of the first and second physical links,

[25.2] to apply the hashing function to the at least one of the frame attributes to produce a hashing key,

[25.3] to calculate a modulo of a division operation of the hashing key by the hashing size, and

[25.4] to select the first and second physical links responsively to the modulo.

[26.1] 26. The apparatus according to claim 25, wherein the control module is arranged to select the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo.

[27.1] 27. The apparatus according to claim 17, wherein the at least one of the frame attributes comprises at least one of a layer 2 header field, a layer 3 header field, a layer 4 header field, a source Internet Protocol (IP) address, a destination IP address, a source medium access control (MAC) address, a destination MAC address, a source Transmission Control Protocol (TCP) port and a destination TCP port.

[28.0] 28. Apparatus for connecting a network node with a communication network, comprising:

[28.1] one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network;

[28.2] a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;

[28.3] a second group of second physical links arranged in parallel so as to couple the one or more interface modules to the communication network; and

[28.4] a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame,

[28.5] the communication network being arranged to provide a communication service to the network node, the service having specified bandwidth requirements comprising at least one of a committed information rate (CR), a peak information rate (PIR) and an excess information rate (EIR), and

[28.6] the first and second groups of physical links being dimensioned to provide an allocated bandwidth for the communication service responsively to the

bandwidth requirements.

[29.0] 29. Apparatus for connecting user ports to a communication network, comprising:

[29.1] one or more user interface modules coupled to the user ports, which are arranged to process data frames having frame attributes sent between the user ports and the communication network,

[29.2] at least one of said user interface modules being bi-directional and operative to communicate in both an upstream direction and a downstream direction;

[29.3] a backplane having the one or more user interface modules coupled thereto and comprising a plurality of backplane traces arranged in parallel so as to transfer the data frames between the one or more user interface modules and the communication network,

[29.4] at least one of said backplane traces being bi-directional and operative to communicate in both said upstream direction and said downstream direction; and

[29.5] a control module, which is arranged to select, for each data frame, responsively to at least one of the frame attributes, a backplane trace from the plurality of backplane traces over which to send the data frame.

[30.0] 30. Apparatus for connecting user ports to a communication network, comprising:

[30.1] one or more user interface modules coupled to the user ports,

[30.2] which are arranged to process data frames having frame attributes sent between the user ports and the communication network;

[30.3] a backplane having the one or more user interface modules coupled thereto and comprising a plurality of backplane traces arranged in parallel so as to transfer the data frames between the one or more user interface modules and the communication network;

[30.4] a control module, which is arranged to select, for each data frame, responsively to at least one of the frame attributes, a backplane trace from the plurality of backplane traces over which to send the data frame;

[30.5] at least some of the backplane traces are aggregated into an Ethernet link aggregation (LAG) group.

[31.1] 31. The apparatus according to claim 29, wherein the control module is arranged to apply a hashing function to the at least one of the frame attributes so as to select the backplane trace.

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioner hereby certifies, in accordance with and in 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 13,987. 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 9, 2023

/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 9, 2023
Manner of service	PRIORITY EXPRESS MAIL
Documents served	Petition for *Inter Partes* Review Under 35 U.S.C. § 312 and 37 C.F.R. § 42.104 of U.S. 7,545,740; Petitioner's Exhibit List; All Exhibits; Petitioner's Power of Attorney.
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

EXHIBIT 2

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner

IPR2023-00402
U.S. Patent No. 8,830,821

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 6

I. INTRODUCTION 9

II. GROUNDS FOR STANDING 9

III. NOTE 9

IV. SUMMARY OF THE '821 PATENT 10

A. Overview of the '821 patent 10

B. Prosecution History of the '821 patent 11

V. LEVEL OF ORDINARY SKILL IN THE ART 12

VI. CLAIM CONSTRUCTION 12

A. "*entity*" 13

VII. RELIEF REQUESTED AND THE REASONS FOR THE REQUESTED RELIEF 13

VIII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE 14

A. Challenged Claims and Statutory Grounds for Challenge 14

B. Ground 1: Claims 1-7, 9-11, 13, and 17-20 are obvious under 35 U.S.C. § 103(a) over Doshi in view of Guichard. 17

1. Summary of Doshi 17

2. Summary of Guichard 20

3. Reasons to Combine Doshi and Guichard 20

4. Claim 1 23

5. Claim 2 46

6. Claim 3 47

7. Claim 4 47

8. Claim 5 50

9. Claim 6 52

10. Claim 7 54

11. Claim 9 54

12. Claim 10 58

13. Claim 11 59

14. Claim 13 59

15. Claim 17 63

16. Claim 18 65

17. Claim 19 65

18. Claim 20 65

C. Ground 2: Claims 8 and 12 are obvious under 35 U.S.C. § 103(a) over Doshi in view of Guichard and Huang 66

1. Summary of Huang 66

2. Reasons to Combine Doshi and Huang 67

3. Claim 8 69

4. Claim 12 71

D. Ground 3: Claims 14-16 are obvious under 35 U.S.C. § 103(a) over Doshi in view of Guichard and Xu. 71

1. Summary of Xu 71

2. Reasons to Combine Doshi and Xu 72

3. Claim 14 75

4. Claim 15 84

5. Claim 16 84

IX. DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE 85

A. Discretionary denial under the *Fintiv* factors is not appropriate 85

1. No evidence regarding a stay 85

2. Parallel proceeding trial date 85

3. Investment in the parallel proceeding 86

4. Overlapping issues with the parallel proceeding 87

5. Identity of parties 88

6. Other circumstances 88

B. Discretionary denial under 35 U.S.C. § 325(d) is not appropriate 88

C. Discretionary denial under *General Plastic* is not appropriate 88

X. CONCLUSION 89

XI. MANDATORY NOTICES 90

A. Real Party-in-Interest 90

B. Related Matters 90

C. Lead and Back-up Counsel and Service Information 90

XII. CLAIMS APPENDIX 92

CERTIFICATE OF WORD COUNT 97

CERTIFICATE OF SERVICE 98

PETITIONER'S EXHIBIT LIST

Ex.1001	U.S. Patent No. 8,830,821 to Cohn et al.
Ex.1002	Prosecution History of U.S. 8,830,821
Ex.1003	Declaration of Dr. Henry Houh under 37 C.F.R. § 1.68
Ex.1004	*Curriculum Vitae* of Dr. Houh
Ex.1005	U.S. Patent Publication No. 2004/0205239 to Doshi et al. ("Doshi")
Ex.1006	"Definitive MPLS Network Designs," by Guichard et al., Cisco Press, 2005 ("Guichard")
Ex.1007	U.S. Patent Publication No. 2003/0117950 to Huang ("Huang")
Ex.1008	"MATE: MPLS Adaptive Traffic Engineering" by Elwalid et al., IEEE, 2001
Ex.1009	"Traffic Engineering with MPLS," by Osborne et al., Cisco Press, 2003
Ex.1010	"Fast Network Re-optimization Schemes for MPLS and Optical Networks" by Bhatia et al., 2006
Ex.1011	RFC3209
Ex.1012	Complaint, *Orckit Corporation v. Cisco Systems, Inc.,* Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Jul. 22, 2022)
Ex.1013	Federal Court Statistics
Ex.1014	Proposed Docket Control Order, *Orckit Corporation v. Cisco Systems, Inc.,* Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Dec. 2, 2022)
Ex.1015	Infringement Contentions, *Orckit Corporation v. Cisco Systems,*

	Inc., Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Nov. 3, 2022)
Ex.1016	U.S. Patent Publication No. 2008/0155257
Ex.1017	U.S. Patent No. 7,624,142
Ex.1018	U.S. Patent Publication No. 2006/0015781
Ex.1019	Definitive MPLS Network Designs Guichard, Copyright Public Record(s)
Ex.1020	Definitive MPLS Network Designs Guichard, Library of Congress Record(s)
Ex.1021	Internet Archive Affidavit, Definitive MPLS Network Designs (Guichard)
Ex.1022	RFC3272
Ex.1023	"Multiprotocol Label Switching (MPLS) Traffic Engineering," Cisco IOS Release 12.0(5)
Ex.1024	IETF "Multiple Metrics for Traffic Engineering with IS-IS and OSPF," Fedyk et al.
Ex.1025	U.S. Patent Publication No. 2011/0141877 to Xu et al., ("Xu")
Ex.1026	U.S. Patent No. 5,699,403 to Ronnen ("Ronnen")
Ex.1027	RFC2026
Ex.1028	U.S. Patent No. 8,102,774
Ex.1029	U.S. Patent Publication No. 2007/0201513
Ex.1030	RFC3469
Ex.1031	U.S. Patent No. 6,980,906

Ex.1032	U.S. Patent No. 7,502,884
Ex.1033	U.S. Patent No. 7,197,418
Ex.1034	U.S. Patent Publication No. 2006/0250948
Ex.1035	U.S. Patent Publication No. 2007/0047469
Ex.1036	U.S. Patent Publication No. 2006/0291391

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 1-20 (the "Challenged Claims") of U.S. Patent No. 8,830,821 ("'821 patent," Ex.1001).

As shown below and confirmed in the Declaration of Dr. Houh (Ex.1003), the features recited in this patent were already known and would have been obvious to a POSITA. *See generally* Ex.1003. The references presented in this Petition render obvious the Challenged Claims, which should be canceled for unpatentability.

II. GROUNDS FOR STANDING

Petitioner certifies that the '821 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.

IV. SUMMARY OF THE '821 PATENT

A. Overview of the '821 patent

The '821 patent generally describes a "system and a method for selecting [Multiprotocol Label Switching] MPLS network transport entities between a first and a second endpoint." Ex.1001, Abstract. According to the '821 patent, the noted system and method provide "protection against multiple span or node failures" by utilizing transport entity pairs that are "fully or partially resource disjoint." Ex.1001, 3:15-18. "[A]n overall cost for each pair of the plurality of transport entities is determined" and a pair, comprising a "working entity" and a "protection entity," is "selected to minimize an overall cost." Ex.1001, 4:13:16, 29:33. For communication, "[a]n active entity may be selected from the working entity and the protection entity." Ex.1001, 4:57-58; Ex.1003, ¶34.

The '821 patent also discloses that an "event" triggers "entity pair reselection." Ex.1001, 4:34-35. The reselection may result in selecting new entity pair or selecting the same pair. Ex.1001, 4:63-5:8. The event trigger may be: a network operational status change, adding an entity, removing an entity, or cost changes for an entity. Ex.1001, 4:35-41; Ex.1003, ¶35.

Figures 2 and 3, reproduced below illustrate flow charts of the '821 patent's methods.

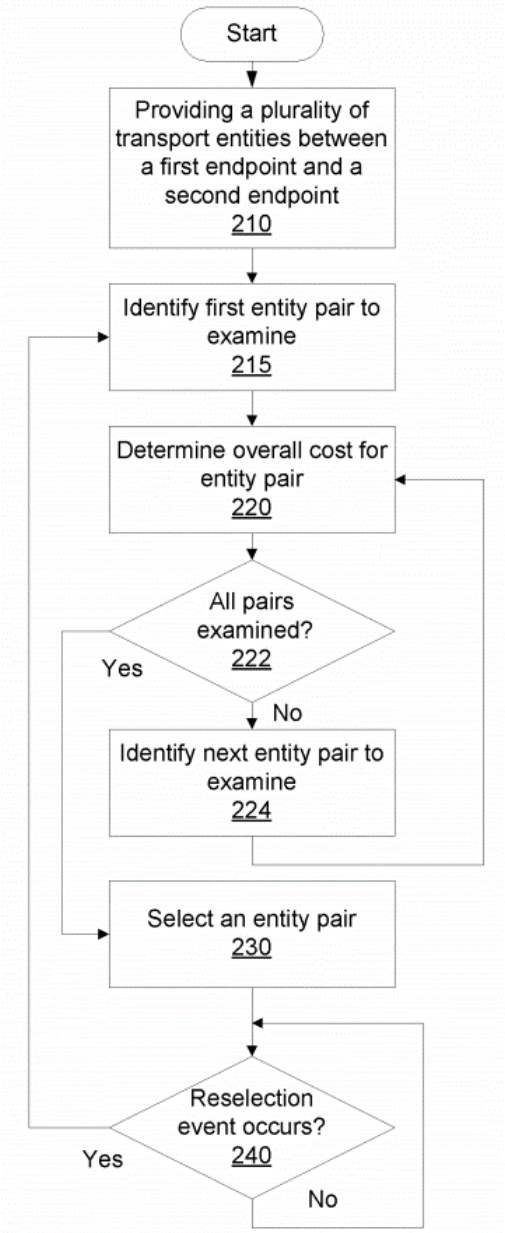


FIG. 2

Ex.1001, Fig. 2.

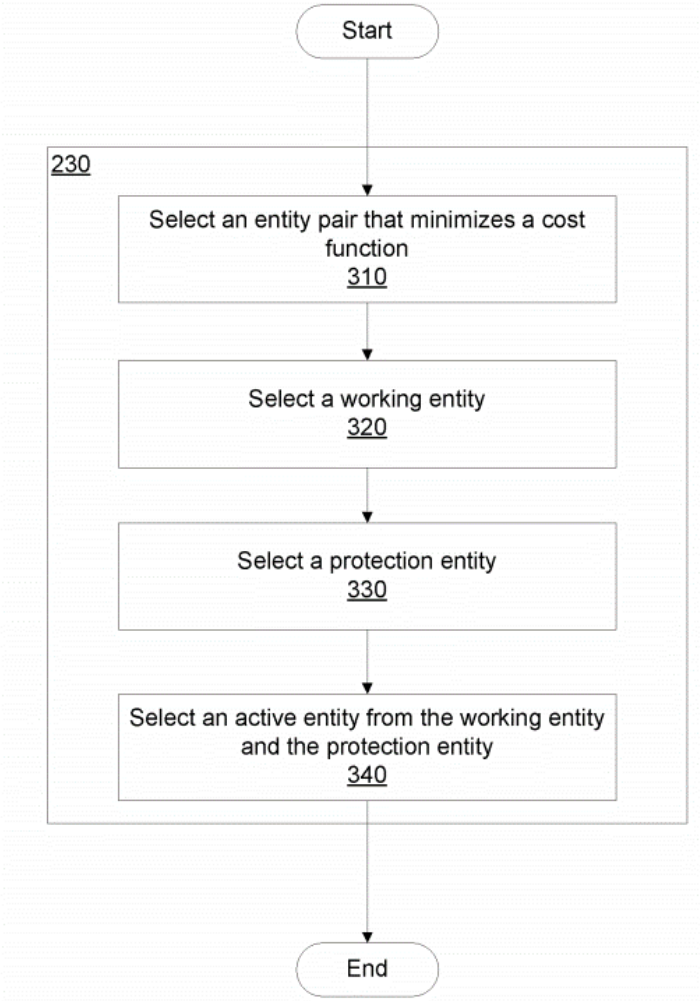


FIG. 3

Ex.1001, Fig. 3.

B. Prosecution History of the '821 patent

The '821 patent claims priority to Provisional application No. 61/499,943 filed on June 22, 2011.

In response to an Office Action, the Applicant amended independent claim 1 "to clarify" the "reselecting said entity pair" limitation, as shown below:

> if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,

Ex.1002, 55, 64. Applicants then argued that the prior art does not teach the amended claim (among other features) because the asserted art only addresses failovers and not reselecting the path pair apart from failover. Ex.1002, 63. The patent issued on September 9, 2014. Ex.1002, 30; Ex.1003, ¶¶38-40.

V. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill in The Art ("POSITA") on June 22, 2011, would have had a working knowledge of network communications and multiprotocol label switching ("MPLS") techniques available at the time. A POSITA would have had a bachelor's degree in computer science, computer engineering, electrical engineering, or equivalent training, and approximately two years of experience working in the field of network communications and been knowledgeable regarding MPLS techniques. Lack of professional experience can be remedied by additional education, and vice versa. Ex.1003, ¶¶27-29.

VI. 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); *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (*en banc*). 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). For clarity, however, Petitioner notes below example embodiments in the specification. Ex.1003, ¶¶41-43.

A. "*entity*"

Each independent claim of the '821 patent recites an "*entity*." The '821 patent discloses that exemplary "MPLS-TP transport entities" include "label switch paths (LSPs) or pseudo wires (PWs)." Ex.1001, 3:22-29. Accordingly, the term "*entity*" includes an MPLS LSP or a PW. Ex.1003, ¶¶44-45.

VII. 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. Petitioner challenges all claims of the '821 patent because they are all asserted in co-pending litigation. Finding the Challenged Claims unpatentable here will reduce the time and expense of '821 patent litigation for all parties.

VIII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

A. Challenged Claims and Statutory Grounds for Challenge[1]

Grounds	Claims	Basis
#1	1-7, 9-11, 13, and 17-20	35 U.S.C. § 103 (Pre-AIA) over Doshi and Guichard
#2	8 and 12	35 U.S.C. § 103 (Pre-AIA) over Doshi, Guichard, and Huang
#3	14-16	35 U.S.C. § 103 (Pre-AIA) over Doshi, Guichard, and Xu

U.S. Patent Publication No. 2004/0205239 to Doshi ("Doshi," Ex.1005) published on October 14, 2004.

The textbook "Definitive MPLS Network Designs," by Guichard et al. ("Guichard," Ex.1006) bears a copyright date of 2005, has an ISBN Number 1-58705-186-9, and states that it was published by Cisco Press in 2005. Ex.1006, 2; Ex.1021. In addition, Guichard was registered as a publication with the Copyright

[1] For each combination presented herein, Petitioner relies on the teachings, and not on a physical incorporation of elements. *See In re Mouttet*, 686 F.3d 1322, 1332 (Fed. Cir. 2012); *In re Etter*, 756 F.2d 852, 859 (Fed. Cir. 1985); Ex.1003, ¶¶66, 213, 233.

Office in 2005, and the Library of Congress online catalog confirms that Guichard was published in 2005. Ex.1006, 2; Ex.1019; Ex.1020. Also, the Internet Archive, which is a well-known archiving website, captured on August 30, 2009 the website Amazon.com, showing that Guichard was available for purchase online and that two reviews were posted in 2005 and one review was posted in 2006, which further support public availability as of those dates. Ex.1021, 1-5. The totality of the evidence demonstrates that Guichard was publicly available to interested persons exercising reasonable diligence, and, therefore, was a printed publication as of 2005. *See Hulu, LLC v. Sound View Innovations, LLC*, IPR2018-01039, Paper 29, at 17-18 (P.T.A.B. Dec. 20, 2018) (precedential); *CIM Maint. Inc. v. P&RO Sols. Grp., Inc*., IPR2017-00516, Paper 8, 18–20 (P.T.A.B. June 22, 2017) (Amazon.com reviews show public availability); *Ex parte Ghalili*, Appeal No. 2020-001741 (P.T.A.B. August 3, 2020) (affirming reference as prior art that was web-archived prior to the effective filing date of the application).

U.S. Patent Publication No. 2003/0117950 to Huang ("Huang," Ex.1007) published June 26, 2003.

U.S. Patent Publication No. 2011/0141877 to Xu ("Xu," Ex.1025) was filed December 15, 2009, and issued in 2013.

Doshi, Guichard, and Huang are each prior art under 35 U.S.C. § 102(b). Xu

is prior art under 35 U.S.C. § 102(e).[2]

Petitioner also appropriately cites additional prior art as 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); *see also K/S HIMPP v. Hear-Wear Techs., LLC*, 751 F.3d 1362, 1365-66 (Fed. Cir. 2014); *Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1363 (Fed. Cir. 2016). For instance, Dr. Houh and this Petition cite to several "Request for Comments" (RFC) documents published by the Internet Engineering Task Force (IETF). To the extent the Board determines that these IETF documents must qualify as prior art "printed publications" for the purposes for which they are cited, the documents do so qualify. *See, e.g.*, Ex.1027, 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);

[2] The '821 patent was prosecuted as a pre-AIA application. Ex.1002, 34. Doshi, Guichard, Huang, and Xu would also be prior art under post-AIA 35 U.S.C. § 102(a).

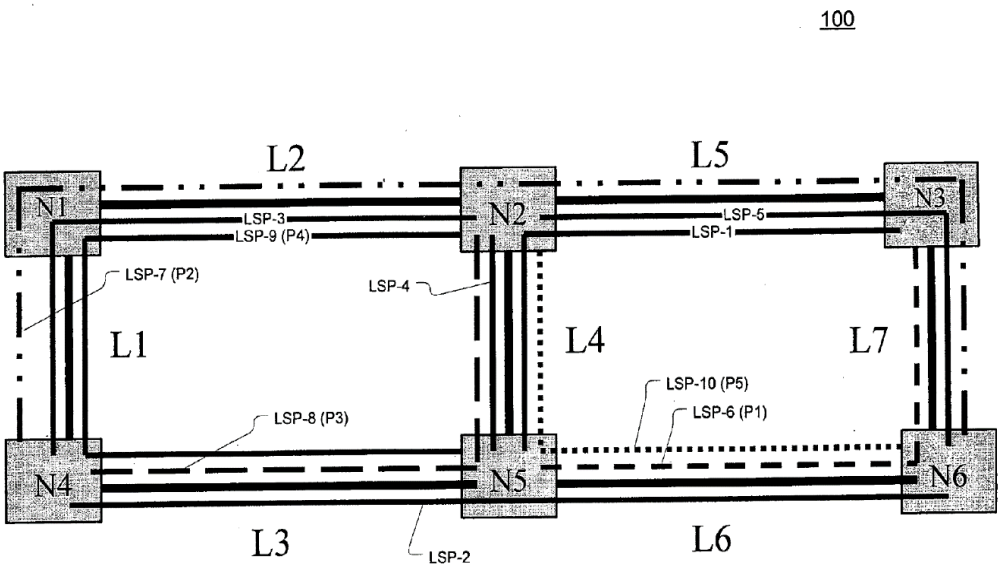
Ex.1003, ¶¶46-47. The Board has repeatedly found IETF documents, including RFCs, to be "printed publications." *See, e.g., Apple, Inc. v. VirnetX, Inc.*, IPR2017-00337, Paper 31, 46-47 (May 30, 2018) (RFCs are "precisely the type of documents whose main purpose is for public disclosure"); *Riot Games, Inc. v. Paltalk Holdings, Inc.*, IPR2018-00130, Paper 11, 30-33 (May 15, 2018) (RFCs are printed publications).

Exhibits 1011, 1022, and 1030 are documents published by the IETF describing MPLS networking standards and improvements, and a POSITA would have reasonably relied upon such documents. Ex.1003, ¶47. In fact, POSITAs frequently cited these exact documents before the '821 priority date. Ex.1005, [0056], [0265] (citing RFC3209); Ex.1035, [0013], [0040]-[0041] (citing RFC3209); Ex.1036, [0011], [0040]-[0041] (citing RFC3209); Ex.1028, 2:22-24 (citing RFC3469); Ex.1029, [0047] (citing RFC3272). Ex.1003, ¶¶47-48.

B. Ground 1: Claims 1-7, 9-11, 13, and 17-20 are obvious under 35 U.S.C. § 103(a) over Doshi in view of Guichard.

1. Summary of Doshi

Like the '821 patent, Doshi (Ex.1005) generally addresses path restoration and recovery. Ex.1005, [0008], Title, Abstract. Doshi, in the context of Figure 1, discloses a shared mesh data network "SMDN 100" that supports an MPLS architecture. Ex.1005, [0048]-[0049]; Ex.1003, ¶¶49-51.



Ex.1005, Fig. 1.

In the above figure, the MPLS architecture includes a plurality of Label Switched Path ("LSP") pairs (e.g., LSP-1 and LSP-6) between two nodes (e.g., between N3 and N5). The LSP pair includes a "primary" (or a "working") LSP and a "protection" (or "restoration") LSP. Ex.1005, [0046], [0052]-[0054], [0076]. The protection LSPs serve as a backup path if the primary LSP fails. Ex.1005, [0052]-[0054], [0266]-[0267]; Ex.1003, ¶52.

Doshi describes a selection method performed by a "network manager," which determines the cost of each LSP pair and then selects the LSP pair with the overall-minimal cost. Ex.1005, [0014], [0142]-[0145], Figs. 10 and 11. The determination considers several factors, including "link utilization, utilization threshold, administrative weight, and sharing degree" and path disjointedness.

Ex.1005, [0142]-[0144]. Doshi explains that the benefit of disjoint paths is that "a failure affecting one of them will not affect the other." Ex.1005, [0040]; Ex.1003, ¶¶50-53.

Doshi's Figure 11, reproduced below, illustrates a logical flowchart for the selection method. Ex.1003, ¶¶52-54.

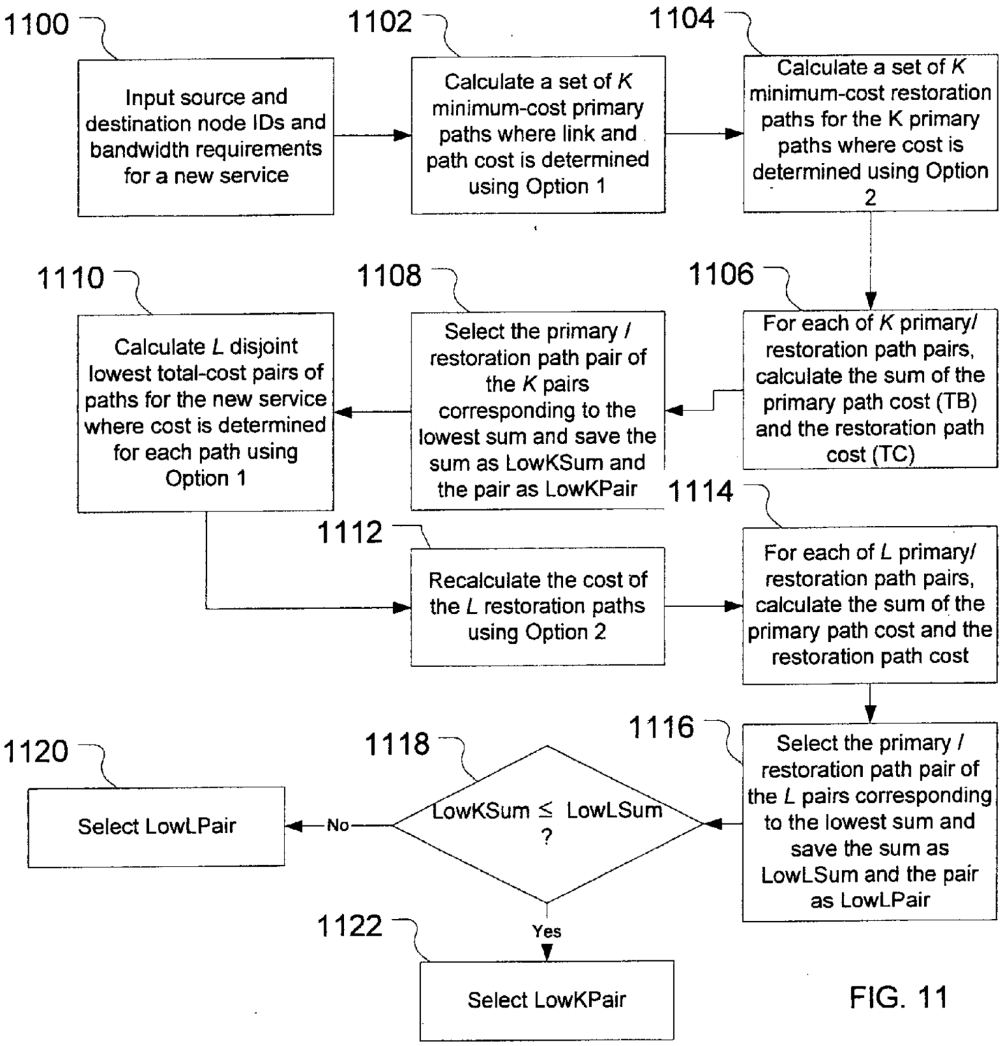


FIG. 11

Ex.1005, Fig. 11.

2. Summary of Guichard

Guichard (Ex.1006) similarly addresses MPLS LSP selection (or reselection), in the context of reoptimization. Ex.1006, 79-81. Guichard teaches that the network may change over time, including nodes and links failing and being restored and new LSPs being added or removed. Ex.1006, 79. These changes present the opportunity to evaluate the available paths and, if warranted, reoptimize the selected LSP by selecting a different LSP if a better path is available. Ex.1006, 79. Guichard provides several reoptimization schemes, including: manual reoptimization, timer-based reoptimization, and event-driven reoptimization. Ex.1006, 80-81; Ex.1003, ¶¶55-56.

3. Reasons to Combine Doshi and Guichard

A POSITA would have been motivated to combine the teachings of Doshi and Guichard. A POSITA would have found it obvious after selecting MPLS LSPs for communication between two nodes, as described by Doshi, to evaluate if the selection should be reoptimized such that a different LSP would be selected when the network changes. Guichard teaches such reoptimization of MPLS LSPs, which it characterizes as "highly desirable." Ex.1003, ¶57.

First, Doshi and Guichard are analogous art to the '821 patent since they generally pertain to MPLS networks and address the problem of selecting LSPs. Ex.1005, [0014], [0049]; Ex.1006, 79-81. Thus, given the similar subject matter, a

POSITA considering Doshi would have naturally considered the teachings of Guichard. Ex.1003, ¶58.

Second, a POSITA would have been motivated to combine the teachings of Doshi and Guichard to produce numerous predictable and beneficial results. Ex.1003, ¶59.

Doshi teaches performing cost calculations and selecting an MPLS LSPs pair that has an overall minimum-cost. Ex.1005, [0014], [0049]. Guichard complements Doshi by teaching how to reoptimize an LSP selection. According to Guichard, it is "**highly desirable** to detect the existence of [network changes] and reoptimize a TE LSP along a better path when it becomes available." Ex.1006, 79. Guichard's reoptimization evaluation may be manually triggered, time-based, or event-driven (e.g., change in the network such as LSP set up and tear down, etc.). Ex.1006, 79-81; Ex.1003, ¶60.

A POSITA would have been motivated to perform reoptimization evaluations, as taught by Guichard, when implementing Doshi's pair calculation and selection techniques to identify a potentially better LSP (e.g., for primary or protection LSP paths) with an even lower overall minimum-cost. Doshi's initial LSP path selection identifies and selects the lowest cost pair of paths at a point in time; however, a POSITA would have recognized that it would be advantageous—indeed "highly desirable," as Guichard puts it—to evaluate the potential for

reoptimization to determine whether the selected LSP pair remains the overall minimum-cost pair or whether there is a better pair available for reoptimization. Performing repeated reoptimization evaluations (based on various triggers), as taught by Guichard, furthers Doshi's stated objective of using the optimal path. Ex.1005, [0172] ("…path selection is a powerful tool that can be used…[for] **optimization**."), [0200] ("…making an **optimal** restoration path choice…"). For instance, reoptimization would allow for selecting a shorter path if one becomes available. Doshi describes attempting to identify the shortest primary and protection paths as part of optimization. Ex.1005, [0151], [0165], [0172]; Ex.1003, ¶61.

The combination represents the use of a known technique (Guichard's reoptimization evaluation in response to various triggers) to improve a similar system or device (Doshi's MPLS architecture) in the same way with predictable results (ongoing optimization). The combination represents what would have been common sense to a POSITA—that it is desirable to evaluate whether selected LSPs remain optimal and to reoptimize the selection if a better option is available. Ex.1003, ¶62.

The results would have been predictable and there would have been a reasonable expectation of success since Doshi and Guichard both address the same technology. Also, a POSITA would have had a reasonable expectation of success

in reoptimizing Doshi's LSP pair selection because Guichard provides significant implementation details, including when to trigger reoptimization evaluations and how to "avoid network instabilities." Ex.1006, 81. The reasonable expectation of success is also evidenced by other skilled artisans that had reoptimized similar systems. Ex.1009, 136; Ex.1010, 1; Ex.1023, 34; Ex.1034, [0027]-[0033], claim 19; Ex.1035, [0024]; Ex.1003, ¶¶63-65.

Implementing the combination would be within a POSITA's skillset since Doshi already sets forth the calculation and selection techniques that would be utilized during the reoptimization evaluation. Also, since Doshi is already "monitoring and updating…network topology" (Ex.1005, [0270]-[0281]) event-based reoptimization evaluation per Guichard, would be directly applicable. Doshi itself acknowledges that a POSITA would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Ex.1005, [300]; Ex.1003, ¶¶64-66.

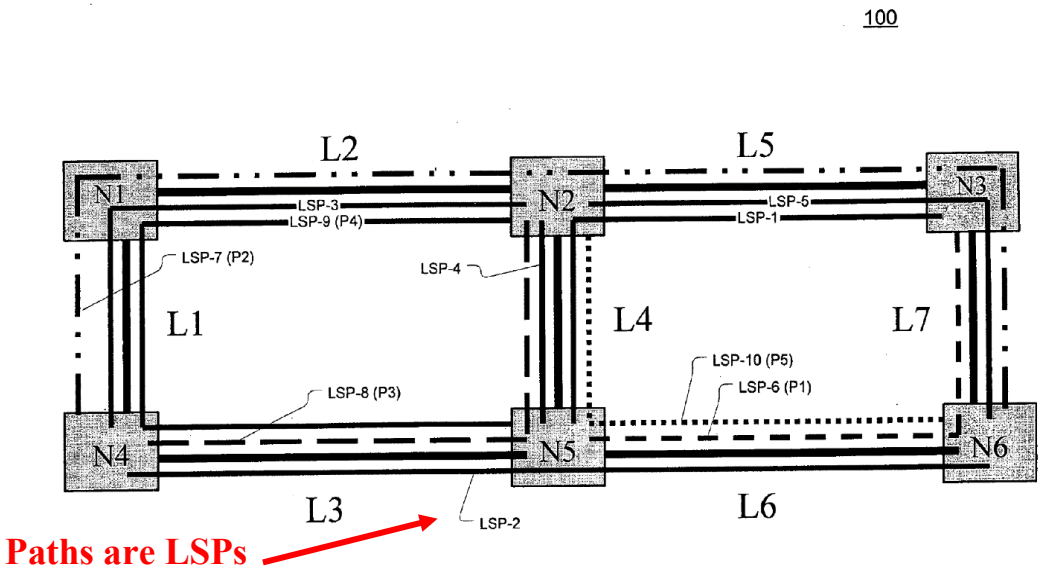
4. Claim 1

[1.0] ***An entity selection method performed by a network device, comprising the steps of:***

As analyzed in the Claim Construction section (§VI.A), the scope of the term "*entity*" includes an MPLS LSP. As shown below, Doshi discloses an MPLS LSP ("*entity*") selection method performed by a network manager ("*network*

device"); Ex.1003, ¶67.

Doshi discloses a shared mesh data network "SMDN 100" that is configured to "support[] the **multiprotocol label switching (MPLS) architecture standard**" of "RFC3031." Ex.1005, [0049]. Doshi's Figure 1, reproduced below, illustrates SMDN 100 that has MPLS LSPs (e.g., LSP-1 to LSP-10) between nodes (e.g., N1-N6). Ex.1005, [0016], [0052]-[0054]; Ex.1003, ¶70.



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶70.

In the MPLS context, Doshi teaches a "**method …implemented by a network manager** for the mesh network." Ex.1005, [0014]; [0049] ("The concepts of the present invention are discussed **in the context of …[a] mesh network (e.g.,… MPLS)**"). The network manager is implemented in a network device, such as a server or distributed partially or fully to nodes. Ex.1005, [0271]; *see also*

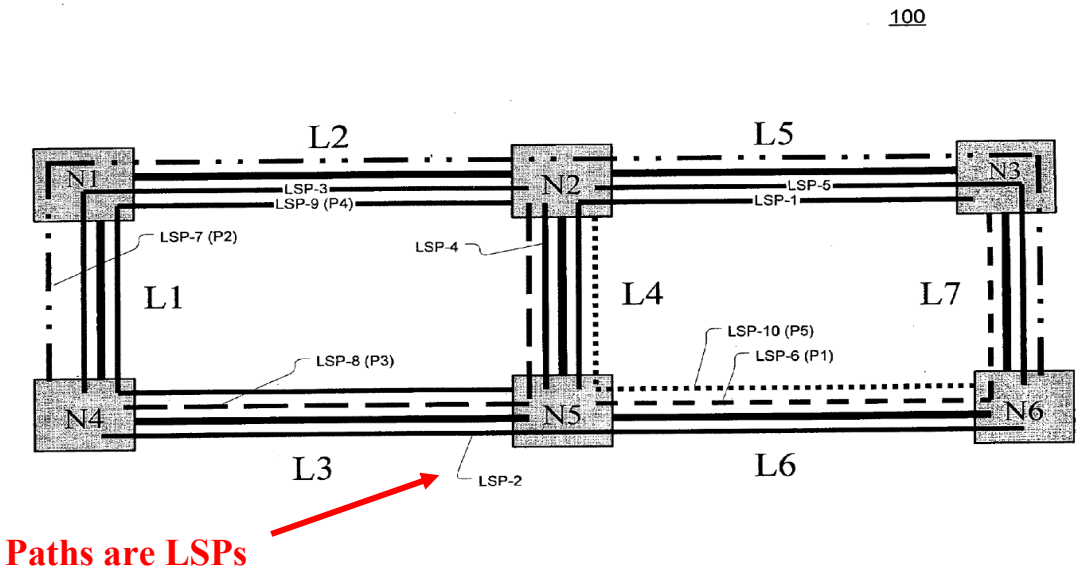
Ex.1005, [0082], claim 11. The network manager performs the method, so that "**primary and restoration paths … are selected** from the plurality of candidate path pairs based on the path cost of each candidate path pair." Ex.1005, [0014]; *see also* Ex.1005, [0142] ("In the first flow, illustrated by steps 1102, 1104, 1106, and 1106, one minimum-cost path pair is **selected** from K candidate pairs by a first method."), claim 11, [0141]-[0145], [0156], [0270]-[0281], Abstract, Figs. 10, 11, 17, 18; Ex.1003, ¶68-69.

A POSITA would have found it obvious, because SMDN 100 supports MPLS, for the "paths" to correspond to LSPs, since those are the paths used for MPLS networks. Ex.1003, ¶71.

Thus, Doshi's description of an MPLS LSP selection method performed by a network manager, renders obvious the preamble. Ex.1003, ¶¶67-73.

[1.1] ***providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;***

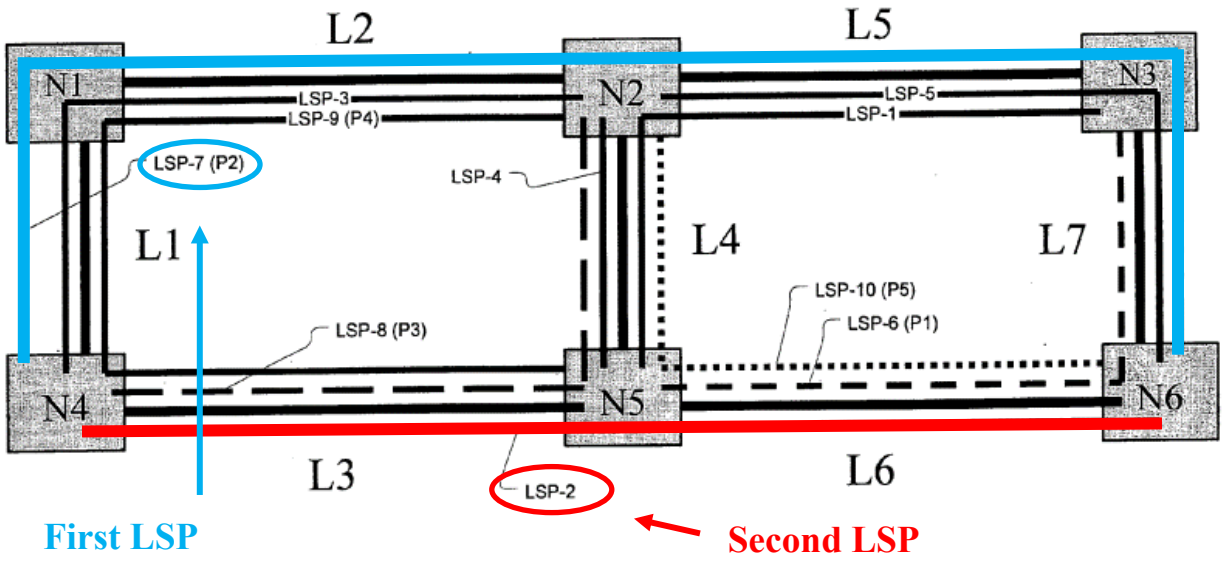
As analyzed above, Doshi's SMDN 100 is configured to support MPLS. Ex.1005, [0049]. Doshi illustrates at Figure 1 that the MPLS network "has been provisioned" with a plurality of LSPs (e.g., "LSP-1" to "LSP-10") between nodes (e.g., N1-N6), at Figure 1. Ex.1005, [0016], [0050]-[0054]; Ex.1003, ¶¶74-76.



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶75.

The LSPs are provisioned "**between two nodes**" that correspond to "**source and destination nodes**" or "**end nodes** of an LSP." Ex.1005, [0050], [0154]-[0156], [0267]; *see also* [0056], [0142], [0188], [0190], [0192], [0195], Fig. 5, Fig. 8, Fig. 11, Fig. 17, Fig. 18, and Fig. 20. Since Doshi's LSPs are provided "*between*" and terminate at two end nodes, a POSITA would have understood that one end node (e.g., source node) corresponds to "*a first endpoint*" and that the other end node (e.g., destination node) corresponds to "*a second end point*." Any two (or more) LSPs provided between two end nodes correspond to "*a plurality of multi protocol label switching (MPLS) transport entities."* Ex.1003, ¶¶77-78.

Below is an example of providing LSP-2 (a "primary" path Ex.1005, [0052]) and LSP-7 (a "protection" path Ex.1005, [0054]) between end node N4 and end node N6. Ex.1003, ¶79.



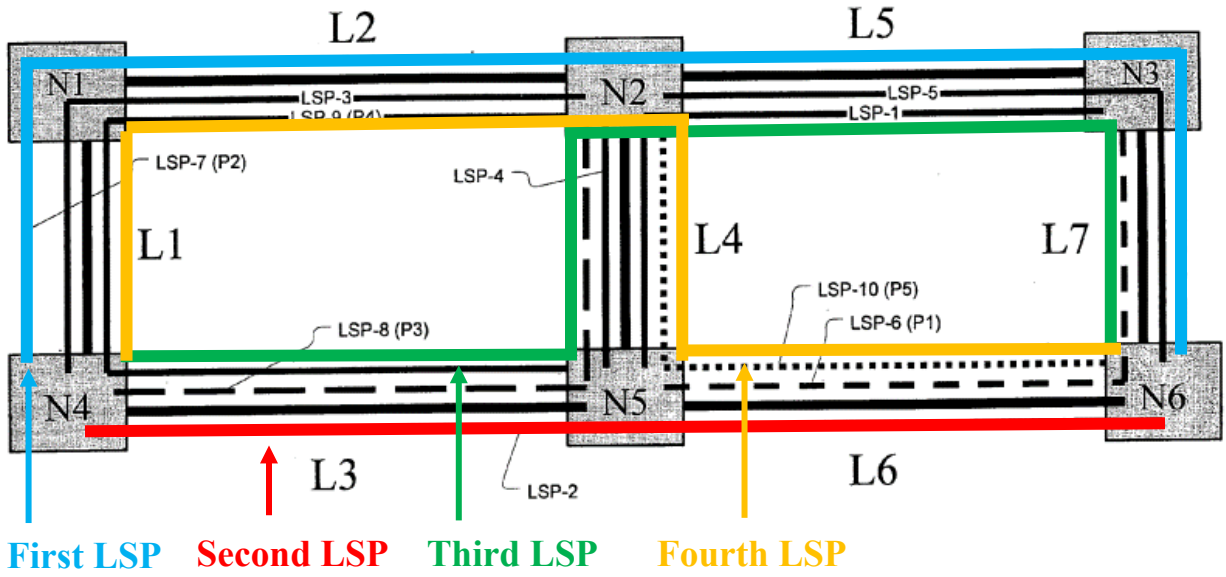
Ex.1005, Fig. 1 (annotated); Ex.1003, ¶79.

Although Figure 1 only illustrates two LSPs between N4 and N6, Doshi makes clear that this figure is merely "**exemplary**" and expressly teaches "**two or more paths between**" the end nodes. Ex.1005, [0016], [0048]-[0055]. Doshi also contemplates making "**[v]arious modifications**" to the exemplary embodiments. Ex.1005, [0300]; Ex.1003, ¶80.

Accordingly, POSITA would have found it obvious to provide additional LSPs between two nodes. Ex.1003, ¶81. That is, it would have been obvious to a person of skill to provide enough LSPs between two nodes to meet bandwidth requirements for communication between the nodes. Doshi itself contemplates more than two LSPs between two nodes by disclosing "two or more paths between" the end nodes and that there are a "plurality of candidate path pairs,"

from which to select for communication. Ex.1005, Abstract, [0014], [0048]. This would mean that any two nodes would have at least two primary LSPs and two protection LSPs, with different primary and protection LSP combinations corresponding to candidate pairs. *See* Ex.1005, [0014], ("…each candidate path pair comprising a candidate primary path and a candidate restoration path…"); *see also* [1.2]-[1.5]; Ex.1003, ¶81

Modified Figure 1 below provides an example of providing two primary LSPs and two protection LSPs between nodes N4 and N6. Ex.1003, ¶82.



Ex.1005, Fig. 1 (modified and annotated); Ex.1003, ¶82.

A POSITA would have recognized that providing more than two LSPs between each node pair (e.g., at least two primary LSPs and two protection LSPs) would be beneficial since doing so facilitates path diversity to reduce the

probability of service interruption between the nodes. Ex.1003, ¶83. For example, the illustrated Third and Fourth LSPs between N4 and N6 offers additional protection against failures along different paths, at least in part. A concurrent failure of, for example, L2 and L6 would cause both the First and Second LSP (the original two LSPs illustrated in Doshi) to go down; however, such a failure would not bring down the Third LSP. Having the Third LSP available would allow traffic to be switched over with minimal interruption of service as compared to if the First and Second LSPs were the only LSPs connecting N4 and N6. Having the Fourth LSP available in the event of concurrent failure of, for example, L3 and L7 (which would affect both the First and Second LSP without affecting the fourth LSP) would similarly allow traffic to be switched over to the Fourth LSP with minimal interruption of service. Ex.1003, ¶83.

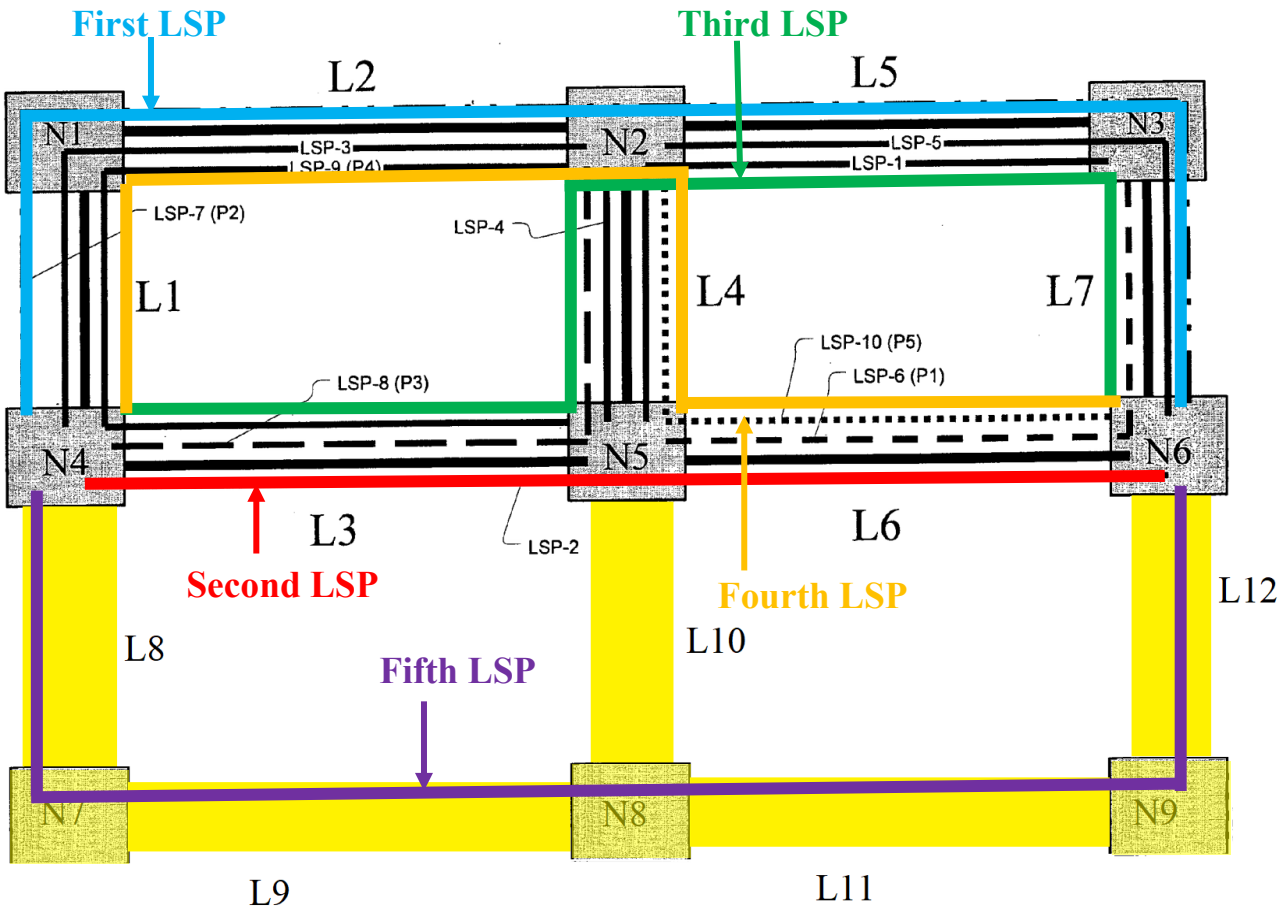
Providing more than two LSPs between two end nodes is desirable because it facilitates reoptimization by providing more potential paths to select from over time. As analyzed below at [1.3]-[1.5], it was known in the art that it is "highly desirable" to reoptimize LSP selection over time. Ex.1006, 79-81. When reoptimizing, a POSITA would have understood to consider all available LSP paths and choose the optimal LSP pair (e.g., primary and protection LSPs). The desirability of optimizing selection is recognized by Doshi, which recognizes the desirability of "optimization." Ex.1005, [0172]; *see also* Ex.1005, [0200]. If only

one primary LSP and one protection LSP was provided between two end nodes, then reoptimization would not be possible in some failure situations. Thus, facilitating reoptimization is another benefit of providing more than two LSPs between each node pair in Doshi. Ex.1003, ¶85.

Another advantage of having is that it allows traffic with different quality of service (QoS) requirements to be allocated and aggregated on different LSPs. *See* Ex.1005, [0265],[0272]. For example, traffic with a high QoS requirement, e.g., video traffic, would be assigned to an LSP configured to handle such high QoS traffic while traffic with a lower QoS requirement would be assigned to a different LSP. Ex.1003, ¶84.

Moreover, it would have been understood that additional nodes with corresponding LSPs not illustrated would be present or may be added over time, as both Doshi and Guichard teach. *See, e.g.,* Ex.1005, [0055] (Describing that other paths and nodes may exist); Ex.1006, 79 ("Network state keeps changing. Links and nodes fail and recover…new TE LSPs are set up…"). A POSITA would have been motivated to provide additional LSPs along new links between new nodes in Doshi, as the network changes, so that the paths between two given nodes are completely disjointed, thereby further reducing the probability of service interruption. *See* Reasons to Combine Doshi and Guichard, § VIII.B.3; Ex.1003, ¶¶85-86.

Modified Figure 1 below illustrates additional nodes N7-N9 with new links L8-L12 that provide corresponding a new Fifth LSP between nodes N4 and N6:



Ex.1005, Fig. 1 (modified and annotated); Ex.1003, ¶87.

As illustrated, the Fifth LSP is completely disjoint with respect to the First and Second LSPs. Ex.1003, ¶¶87-88

Thus, Doshi in combination with Guichard discloses providing a plurality of MPLS LSPs (e.g., at least two primary LSPs and at least two protection LSPs) between a first end node (e.g., source node) and a second end node (e.g., destination node), which renders obvious this limitation. Ex.1003, ¶¶74-89.

[1.2] ***determining an overall cost for each entity pair of said plurality of entities;***

Doshi calculates the cost of each path (e.g., each primary or protection path) as "the sum of the link costs." Ex.1005, [0125], Fig. 10. Then, Doshi's path cost is used to determine the cost of each path pair. "**For each of a plurality of candidate path pairs…a path cost associated with said each candidate path pair is generated**." Ex.1005, [0014]; *see also* Ex.1005, Abstract, [0026], [0046], [0050]-[0054], [0076], [0082], [0130], [0124]-[0181], Fig. 1, Fig. 10, Fig. 11, Fig. 12, and Fig. 13; Ex.1003, ¶¶90-94.

Doshi's cost determination takes into account whether path pair is "disjoint," as well as other factors, such as the "link utilization, utilization threshold, administrative weight, and sharing degree." Ex.1005, [0142], Ex.1005, [0146]-[0156], Fig. 11. Based on these factors, Doshi identifies the path pair with an "**overall-minimum cost**." Ex.1005, [0142]; Ex.1003, ¶¶95-97.

Doshi's Figure 11 (building upon Figure 10's example for determining individual path cost), shows that the cost of each candidate pair is determined by summing a primary path cost and a protection path cost. Ex.1003, ¶95.

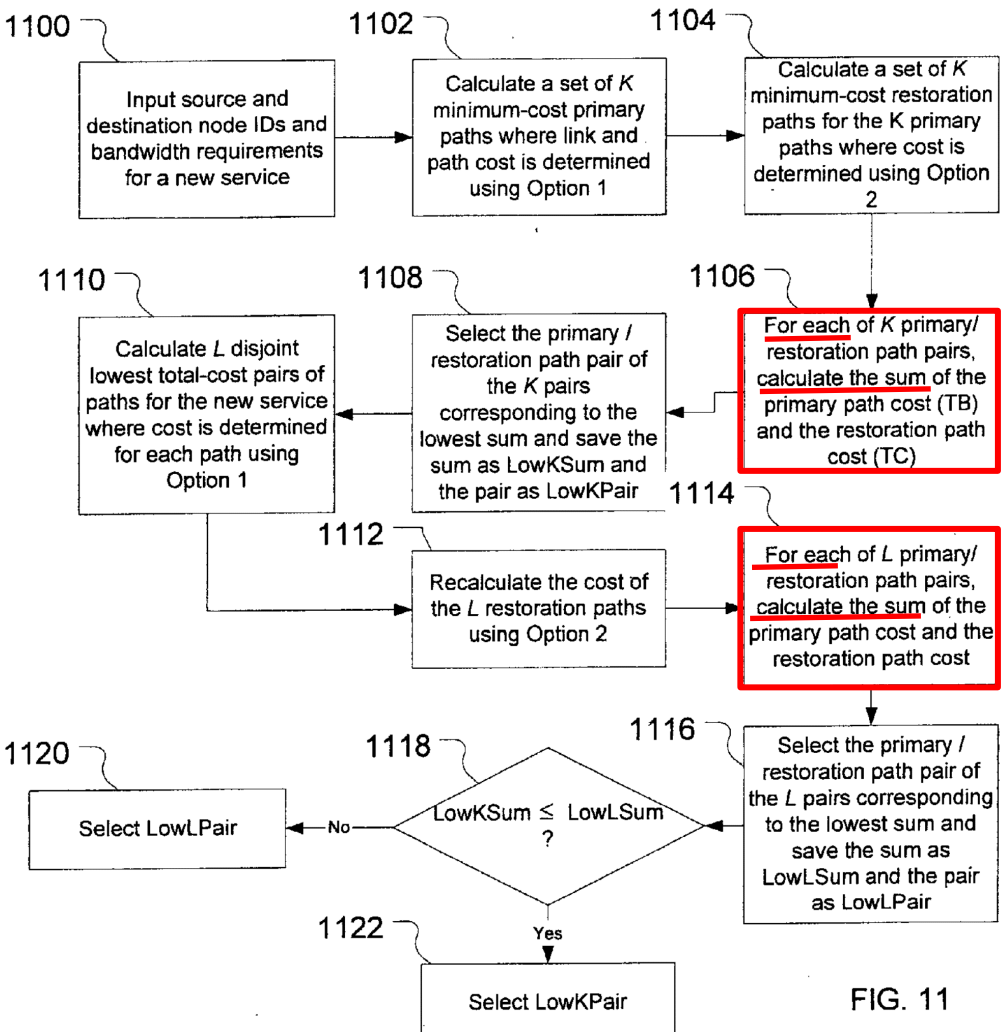


FIG. 11

Ex.1005, Fig. 11 (annotated); Ex.1003, ¶95.

A POSITA would have understood that Doshi's cost determination—which considers a myriad of factors such as link utilization, sharing degree, and path disjointedness—represents the "*overall cost*" for each candidate path pair. Ex.1003, ¶¶96-97.

Consistent with the analysis at [1.0]-[1.1], it would have been obvious for the path to be implemented as LSPs, since Figure 1 supports an MPLS architecture

and provides a plurality of LSPs. Also, a POSITA would have found it obvious for the candidate path pairs to be from the provided plurality of MPLS LSPs (at least two primary and at least two protection LSPs) between the end nodes (See [1.1]), because these LSPs would have already been in existence and available for immediate use for the new service. Ex.1003, ¶98.

By determining the cost for the provided LSPs, assessment can be made whether the provided LSPs have sufficient resources for the new service—without having to allocate additional resources. Considering the provided LSPs for the new service would allow such LSPs to be more fully utilized. Ex.1008, 1306 (describing the known problem of underutilized LSPs and solution of allocating traffic to such LSPs); Ex.1003, ¶99

Consistent with the analysis at [1.3]-[1.5] below, in view of Guichard's teaching of ongoing reoptimization efforts, it would have been an obvious aspect of path selection to perform Doshi's cost calculations repeatedly to attempt to determine whether the original LSP pair still has the overall minimum-cost or if there is another LSP pair that is more optimal. A POSITA would have understood that the cost calculations for such reoptimization efforts would consider existing paths, including, for example, an originally selected LSP pair that is currently being utilized. Ex.1003, ¶¶100-101.

[1.3] ***selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and***

First, Doshi discloses that "**primary and protection paths**…**<u>are selected</u> from the plurality of candidate path pairs based on the path cost of each candidate path pair**." Ex.1005, [0014]; *see also* Ex.1005, Abstract, [0046], [0076], [0142]-[0145], Claim 1, Claim 11. Consistent with the analysis at [1.1]-[1.2], it would have been obvious for the "plurality of candidate path pairs" (pairs of primary and protection LSPs) to correspond to the provided plurality of LSPs of the MPLS architecture. Ex.1003, ¶¶102-103.

Doshi further teaches that the selection is based on the overall cost, since "**the lowest-cost pair…is <u>selected as the overall minimum-cost pair</u>**." Ex.1005, [0142]; *see also* Ex.1005, [0026], [0125]-[0139], [0143]-[0145]. Doshi's Figure 11 and illustrates selecting the overall minimum-cost pair. Ex.1003, ¶104

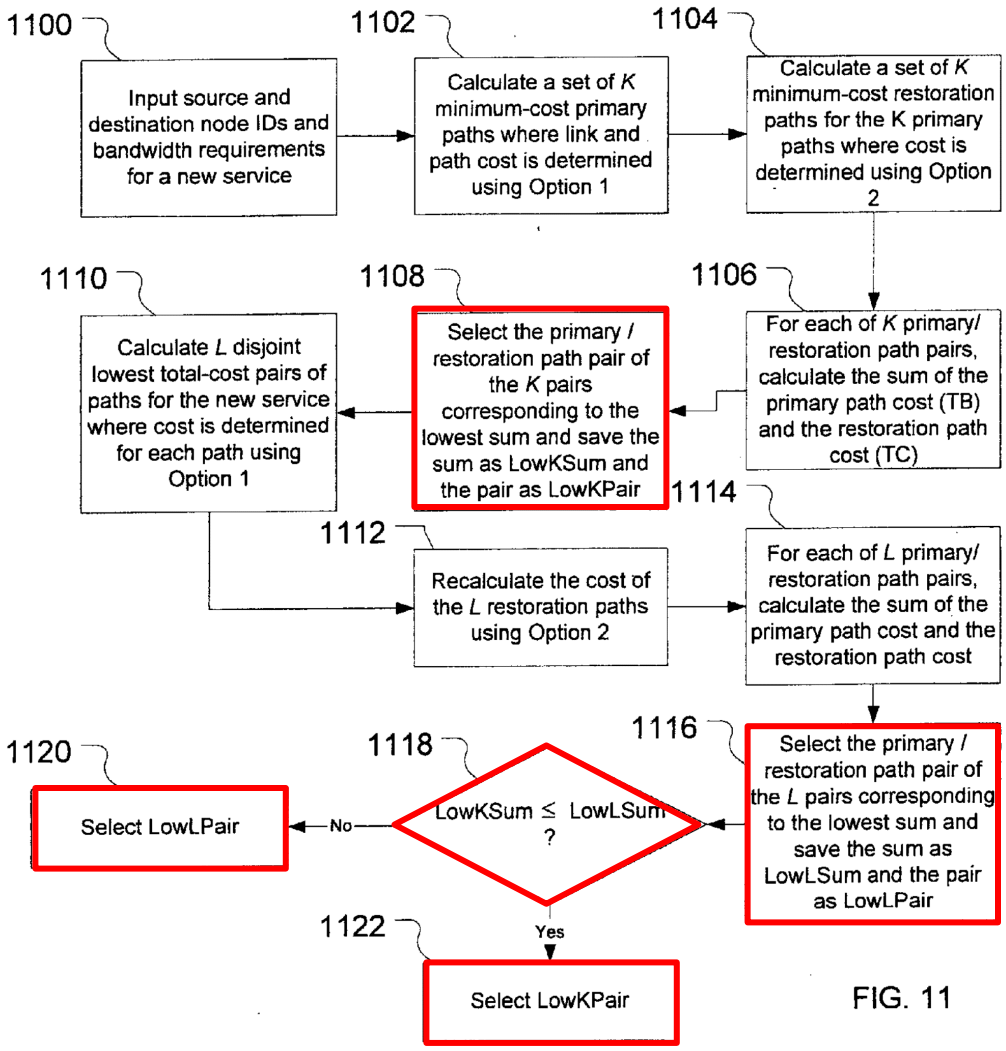


FIG. 11

Ex.1005, Fig. 11 (annotated); Ex.1003, ¶104.

Therefore, Doshi discloses selecting an LSP pair from the plurality of MPLS LSP based on the selected pair having the overall minimum-cost. Ex.1003, ¶¶102-105.

Second, Guichard teaches that it is "highly desirable" evaluate the selected LSP and reoptimize the selection "[i]f a better path is found" or if a "more optimal (shortest) path exists." Ex.1006, 79. Guichard contemplates, for example, "Manual

reoptimization" and "Timer-based reoptimization." Ex.1006, 81; *see also* Ex.1006, 79. Guichard merely describes what would have been common sense to a POSITA—reoptimizing a system to use the best available LSP. In other words, a POSITA would have understood that Guichard teaches that original path selections should be evaluated on an ongoing basis to determine whether better or more optimal paths exist. Ex.1003, ¶¶106-109.

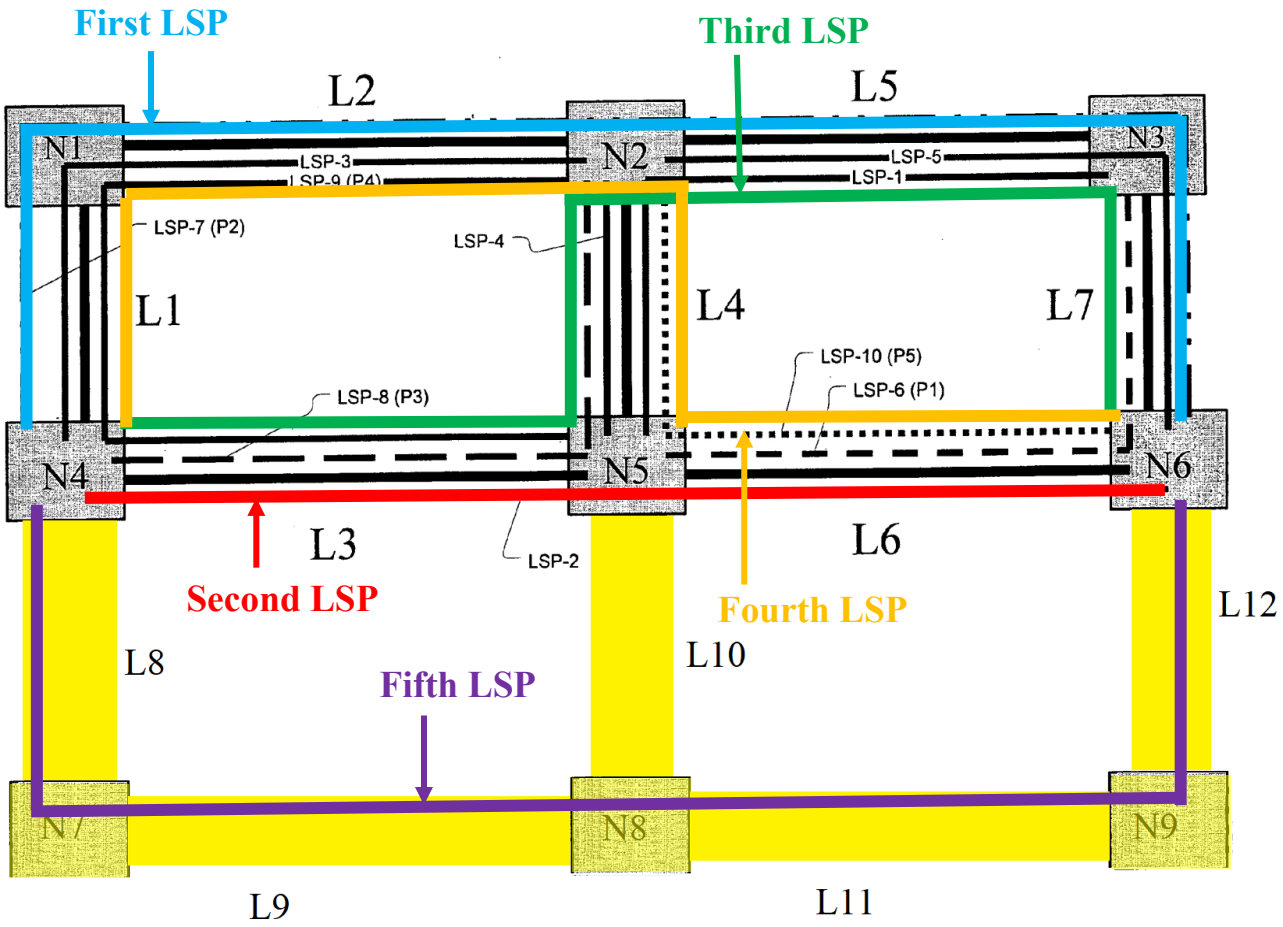
In view of Guichard's teachings, it would have been obvious to reoptimize Doshi's LSP selection by performing cost calculations repeatedly over time (e.g., based on a manual trigger or timer-based) to determine whether the originally selected LSP pair still has the overall minimum-cost or if another LSP pair should be selected. A POSITA would have understood that the cost calculations for such reoptimization efforts would consider existing LSP paths, including, the originally selected LSP pair that had the overall minimum-cost. If the originally selected LSP pair still has the overall-minimum cost, a POSITA would have understood that Doshi's selection techniques identified above would reselect the original pair. On the other hand, if another LSP pair is determined to have the overall-minimum cost, then that LSP pair would be selected. Each such selection is an example of "*selecting,*" as claimed. Ex.1003, ¶110

It would have been obvious to a POSITA, in view of Guichard, to perform ongoing reoptimization attempts, when implementing Doshi's MPLS architecture,

so that an optimum LSP pair (e.g., overall minimum-cost pair) is selected over time for communication between two end nodes. In implementing the combination, a POSITA would have evaluated all existing available LSPs (including previously selected LSPs) by repeating Doshi's candidate pair cost determination (See [1.2]) and selection techniques (discussed above) in response to a manual trigger or expiration of a timer to thereby select the LSP pair with an overall minimum-cost. Ex.1003, ¶110.

A POSITA would have been motivated to perform ongoing reoptimization to identify potentially better primary and protection LSP paths (for the LSP pair between two end nodes) that minimize the overall cost. Ex.1006, 79 ("It is then **highly desirable** to detect…and reoptimize a TE LSP along a **better path** when it becomes available."); Ex.1005, [0172] ("In general, path selection is a powerful tool that can be used to achieve…**optimization**.") [0200] ("…making an **optimal** restoration path choice…"); *see also* Reasons to Combine Doshi and Guichard, § VIII.B.3; Ex.1003, ¶111.

Below is an example of such a selection, per the combined teachings, in the context of Doshi's modified Figure 1. Ex.1003, ¶112.



Ex.1005, Fig. 1 (modified and annotated); Ex.1003, ¶112.

Take, for instance, where an original LSP pair (e.g., First LSP and Second LSP) was being utilized for communications between end nodes N4 and N6. In response to a manual trigger or expiration of a timer, as Guichard teaches, potential reoptimization of the LSPs would take place to determine if there is a better or more optimal LSP. All things being equal, for example, in the circumstance where the links of the First LSP are heavily utilized as compared to the links of the Fifth LSP, it would have been obvious for Doshi's cost determination to result in a new

pair (e.g., the Fifth LSP as the primary and the Second LSP as the protection) as having the overall minimum cost. In such a circumstance, these LSPs would be selected for further communications between end nodes N4 and N6 while the First LSP would be replaced by the Fifth LSP. Ex.1003, ¶¶113-114.

[1.4] ***if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,***

First, Guichard teaches "**Event-driven reoptimization**" where "**it may be desirable to trigger a reoptimization upon the occurrence of a particular event**, such as the restoration of a link in the network." Ex.1006, 81. Additionally, "the evaluation of any potential reoptimization" is triggered "each time a new link is advertised as newly operational." Ex.1006, 81; *see also* Ex.1006, 79-80 (triggering reoptimizing when links and nodes fail or recover and when new LSPs are set up or torn down and a better path is available). Ex.1003, ¶¶115-116.

A POSITA would have understood that Guichard generally teaches that the current LSP path selection should be evaluated following a trigger event (e.g., LSP set up, LSP torn down, restoration of a link, or new link) to determine whether a more optimal LSP path is available. Also, a POSITA would have recognized that this evaluation may or may not result in a reoptimization. *See* Ex.1006, 81 (referring to "evaluation of any **potential** reoptimization."); Ex.1035, [0057]-[0059] (recognizing that if "new cost exceeds the old cost … no optimization of

the TE-LSP is performed."). A person of skill would have understood that there is only a potential for reoptimization since if the evaluation indicates that the currently utilized path is still the best option, then no reoptimization would occur and the previously utilized path is reselected. Ex.1003, ¶117.

It would have been obvious to a POSITA, in view of Guichard, to perform event-driven reoptimization, when implementing Doshi's MPLS architecture, to determine if there are better LSP paths available or if the currently selected LSP pair represents the overall minimal cost. *See* Reasons to Combine Doshi and Guichard, § VIII.B.3; Ex.1003, ¶118.

In implementing the combination with Doshi, a POSITA would have evaluated any restored and new LSPs along with all existing LSPs (including the currently utilized LSP pair) by repeating Doshi's candidate pair cost determination and selection techniques (See [1.2]-[1.3]) in response to a trigger event to potentially reoptimize the selection. Ex.1003, ¶119.

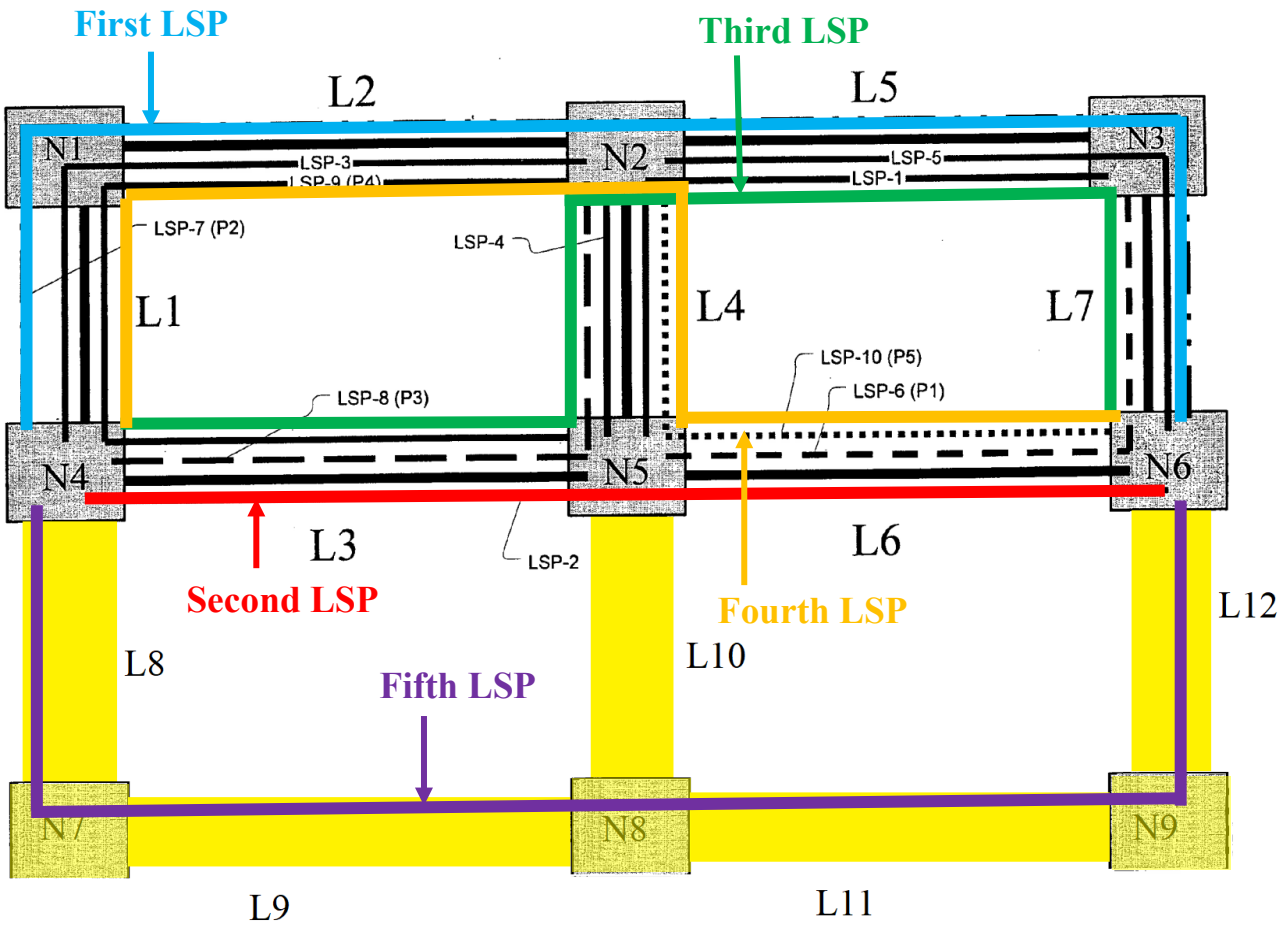
It would have been obvious to a POSITA that repeating Doshi's candidate path pair cost determination and selection would—in certain instances—result in the reselection of the currently utilized LSP pair. This would be the case where the currently utilized LSP pair still provides the overall minimum-cost, despite an LSP path being restored, a new LSP path being set up, or some other change in the network. In the combination, Doshi's currently utilized LSP pair would be

repeatedly evaluated and reselected, per Guichard's event-driven reoptimization technique, if it continues to have the overall-minimum cost per Doshi's cost determination (See [1.2]). Ex.1003, ¶120.

Thus, Doshi in combination with Guichard discloses an event triggering candidate pair cost determination and, if the cost determination indicates that the currently used LSP pair still has the overall-minimum cost, reselecting the LSP pair, which renders obvious "*if an entity pair reselection event occurs, reselecting said entity pair*." *See also* [1.5], below, analyzing further the reselection event. Ex.1003, ¶121.

Second, as discussed above at [1.0]-[1.3], Doshi's plurality of candidate path pairs include a plurality of LSPs (e.g., at least two primary and at least two protection LSPs) between a pair of end nodes. The currently used primary and protection LSP pair would be initially selected based on the overall minimum cost (see [1.3]) and the remaining LSPs would be available (with any restored or new LSPs in the network that generate the trigger event, per Guichard) as potential replacements during reoptimization. Thus, in the proposed combination, the reselection is "*from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair*." Ex.1003, ¶122.

As an example, consider the modified Figure 1, per the combined teachings, discussed at [1.3], and reproduced below. Ex.1003, ¶123



Ex.1005, Fig. 1 (modified and annotated); Ex.1003, ¶123.

As noted at [1.3], the selection step selected the Fifth LSP as the primary path and the Second LSP as the protection path. When performing event-driven based reoptimization, per Guichard, an overall cost for each LSP pair permutation would be determined, per Doshi. Ex.1005, [0014] ("…a path cost associated with

said each candidate path pair is generated…"). If the cost determination indicates that the currently used LSP pair (e.g., Fifth LSP and Second LSP) still has the overall-minimum cost, then the same LSP pair would be reselected. Ex.1005, [0142] ("…the lowest-cost pair…is selected as the overall minimum-cost pair…"); Ex.1003, ¶124.

Since Doshi's cost determination is performed on each candidate path pair, it would have been obvious to a POSITA for the reselection to be from the group consisting of the currently used LSP pair (e.g., Fifth LSP and Second LSP) and a replacement LSP pair comprising at least one LSP distinct from the currently used LSP pair. Ex.1003, ¶125. In the present example, the reselected LSP pair (e.g., Fifth LSP and Second LSP) would correspond to "*said entity pair*" while each remaining permutation of LSP pairs not resected would correspond to "*a replacement entity pair comprising at least one entity distinct from the entities of said entity pair*." Ex.1003, ¶125.

Third, to the extent argued that "*reselecting said entity pair*" requires selecting at least one different LSP as part of the pair instead of reselecting the same two LSPs, the proposed combination also renders that obvious. A POSITA would have found it obvious to select a different LSP pair, when the different LSP pair is determined to have a lower overall-minimum cost than the currently used LSP pair. To illustrate, in the present example, if the First LSP and Second LSP are

now determined to have the lowest overall-minimum cost, then that LSP pair would be reselected. Ex.1003, ¶¶115-127.

[1.5] ***wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.***

Guichard teaches that reoptimization evaluation is triggered by various events ("*entity pair reselection event*"). For example, Guichard's event-driven reoptimization is triggered "if a **link is restored** in the network" and "each time a **new link is advertised as newly operational.**" Ex.1006, 81. A POSITA would have understood that restoring a link restores LSPs of the link, which may provide a more optimal path. As discussed above at [1.0]-[1.4], the plurality of MPLS LSPs between end nodes correspond to "*said plurality of transport entities*." It follows that restoring a link that previously went down to "newly operational," means that the operational status of the link's LSPs has also changed to operational, which discloses "*an operational status change for one of said plurality of transport entities*." Ex.1003, ¶¶128-130.

Another network change event contemplated by Guichard is when "**new…LSPs are set up**" and when "others may be **torn down**." Ex.1006, 79. Because these events have the potential for "a more optimal path [to] appear in the network. It is then highly desirable to detect the existence of such a path and

reoptimize a TE LSP along a better path when it becomes available." Ex.1006, 79. It follows that a new LSP set up between end nodes would be added to the plurality of MPLS LSPs between the end nodes, which corresponds to "*adding an entity to said plurality of transport entities*." Conversely, tearing down an LSP between the end nodes would remove the LSP from the plurality of MPLS LSPs between the end nodes, which corresponds to "*removing an entity from said plurality of transport entities*." Ex.1003, ¶¶132-133.

5. Claim 2

[2.1] ***The method of claim 1, wherein said step of selecting an entity pair further comprises: selecting a working entity from said entity pair; and selecting a protection entity from said entity pair.***

As discussed above at [1.3], Doshi discloses selecting an LSP pair from the plurality of MPLS LSPs. The selection includes selecting a primary path (Doshi also refers to as a "working path") and selecting a "protection path," which corresponds to "*selecting a working entity*" and "*selecting a protection entity*." Ex.1005, [0014], [0046]; *see also* Ex.1005, Abstract, [0142]-[0145], Claim 1, Claim 11; Ex.1003, ¶¶134-135.

The '821 patent describes "switch[ing] between a working entity and a protection entity" after failure. Ex.1001, 1:51-43; 2:35-36. Doshi similarly describes "Switching Between Working and Protection LSPs" "after detecting a failure." Ex.1005, [0266]-[0267]. Doshi's working LSP corresponds to a "*working*

entity" and its protection LSP corresponds to a "*protection entity*." Since Doshi's working and protection LSPs are part of the selected LSP pair, a POSITA would have understood that the selection of the working and protection LSPs are "*from said entity pair*." Ex.1003, ¶¶136-137.

6. Claim 3

[3.1] ***The method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity.***

As discussed at [2.1], Doshi describes selecting a working LSP ("*working entity*") and a protection LSP ("*protection entity*"). Doshi further describes "Switching Between Working and Protection LSPs" and explains that "after detecting a failure, end nodes of an LSP switch traffic from a primary (i.e., working) LSP to its corresponding protection LSP." Ex.1005, [0266]-[0267]. Doshi's working LSP is the path actively carrying traffic absent a failure while Doshi's protection LSP becomes active when traffic is switched over following a failure. Thus, Doshi's switching decisions, which select which of the working LSP and protection LSP is in use, render obvious this limitation. Ex.1003, ¶¶138-140.

7. Claim 4

[4.1] ***The method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.***

As discussed at [1.2]-[1.3], Doshi calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair and sums the path costs

to find the cost of the pair together. The cost calculation considers various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness, and which represent the "*overall cost"* for each candidate path pair. The calculation used to determine the "*overall cost*" corresponds to the claimed "*overall cost function*." Ex.1003, ¶142.

Doshi's Figure 11 and accompanying description teach calculating and selecting an "**overall minimum-cost"** pair. Ex.1005, [0142]; *see also* Ex.1005, [0026], [0125]-[0139], [0143]-[0145]. Ex.1003, ¶¶143-144.

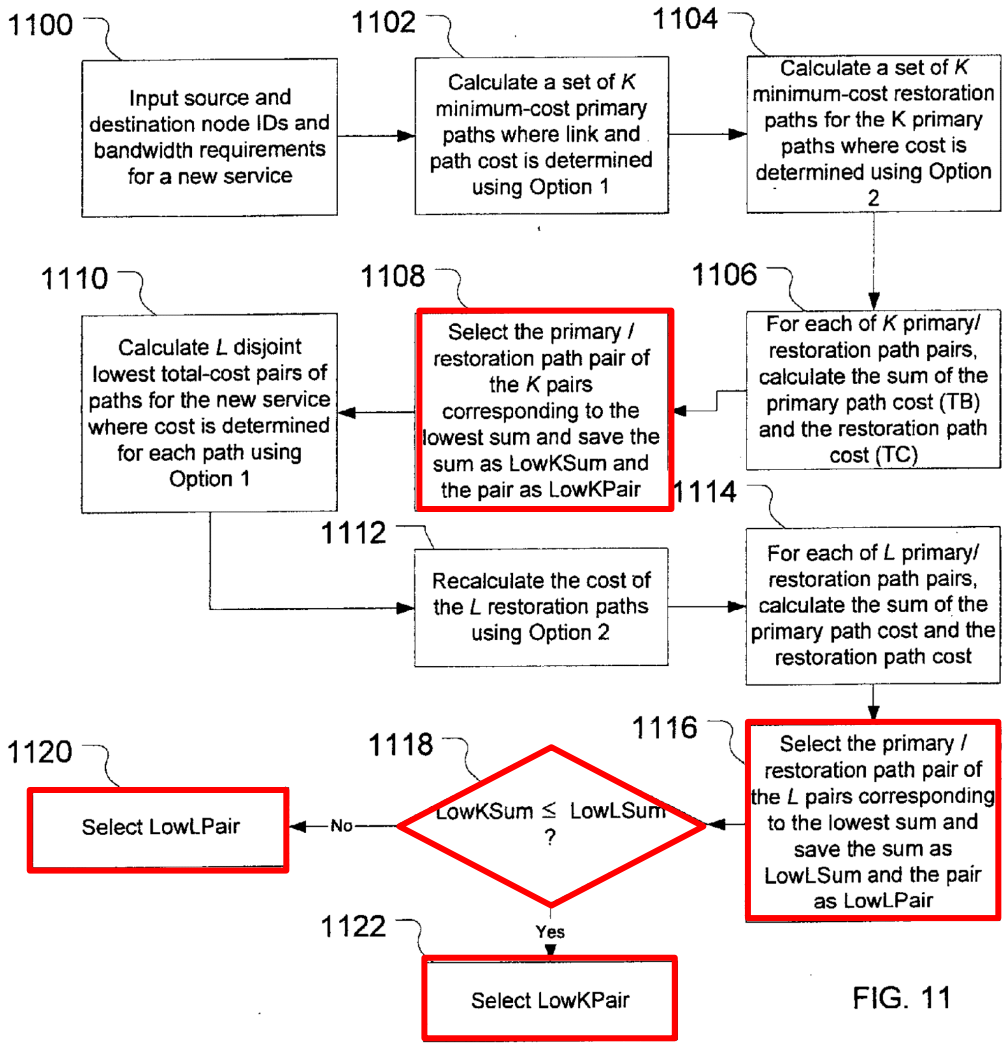


FIG. 11

Ex.1005, Fig. 11 (annotated); Ex.1003, ¶143.

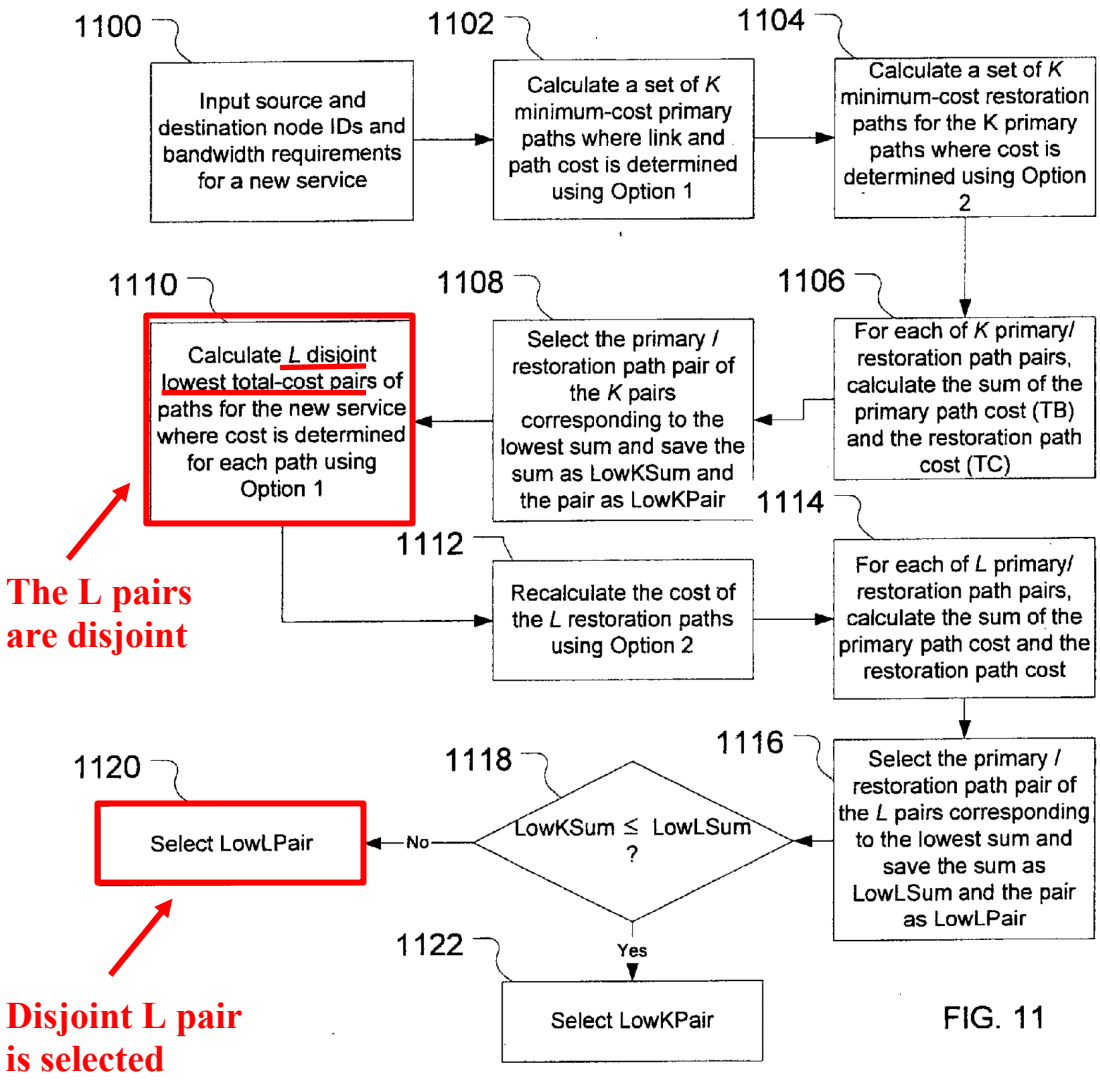
A POSITA would have understood that the selected "overall minimum-cost" pair is the LSP pair having the lowest overall cost function in consideration of various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness. Thus, Doshi renders obvious this limitation. Ex.1003, ¶¶141-145.

8. Claim 5

[5.1] ***The method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.***

As discussed at [1.2]-[1.3], and [4.1], Doshi calculates the path cost of each primary "*working entity*" and protection LSP "*protection entity*" in a candidate path pair and sums the path costs to find the cost of the pair together. The cost calculation considers a myriad of factors such as link utilization, sharing degree, and path disjointedness and which represent the "*overall cost"* for each candidate path pair. The calculation used to determine the "*overall cost*" corresponds to the claimed "*overall cost function*." Ex.1003, ¶147.

Doshi describes selecting primary LSP and protection LSPs that are "**strictly disjoint**" from each other, having "no common links or nodes other than their common ingress and egress nodes." Ex.1005, [0053]; Ex.1003, ¶¶148-150.



Ex.1005, Fig. 11 (annotated); Ex.1003, ¶148.

When two paths are strictly disjoint, "**then a failure affecting one of them will not affect the other**." Ex.1005, [0040]; *see also* Ex.1005, [0155] (Failure "**protection is made possible** by selecting…disjoint [paths]."); Ex.1005, [0056], [0062], [0080], [0131], [0142]-[0156], [0287], Fig. 1. By considering path disjointedness in its cost calculation and selecting a pair of primary and protection

paths that are node and link disjoint such that the failure of one does not affect the other, Doshi renders obvious this limitation. Ex.1003, ¶¶146-151.

9. Claim 6

[6.1] ***The method of claim 4, wherein said overall cost function comprises a predefined entity cost metric.***

As discussed at [1.2]-[1.3] and [4.1], Doshi calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair and sums the path costs to find the cost of the pair together. The cost calculation considers various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness—which represent the "*overall cost"* for each candidate path pair. The calculation used to determine the "*overall cost*" corresponds to the claimed "*overall cost function*." Ex.1003, ¶¶152-153.

Each of these factors, considered in the cost calculation, is predefined in the sense that the cost calculation is pre-configured to consider each factor. For example, before a calculation is performed, it is understood that the calculation will consider link utilization, sharing degree, and path disjointedness. Additionally, the "utilization threshold" (e.g., nominally set to 80%) and "administrative weight" (e.g., a coefficient for scaling) are predefined metrics. Ex.1005, [0128], [0141]. Each of these metrics has an impact on the cost of the path being evaluated. Ex.1005, [0142]. Thus, each of link utilization, utilization threshold, administrative

weight, sharing degree, and disjointedness correspond to "*a predefined entity cost metric*." Ex.1003, ¶154.

Additionally, traffic engineering metrics correspond to "*predefined entity cost metric*" in the '821 patent. Ex.1001, 5:49-51. A POSITA would have understood that Doshi's link utilization, utilization threshold, and administrative weight each correspond to traffic engineering data—it was well known that traffic engineering data includes such metrics. *See* Ex.1022 (using "administrative weights" to implement traffic engineering decisions.); Ex.1023 (setting an administrative weight with an "mpls traffic-eng administrative-weight" command, setting a bandwidth threshold level of reserved bandwidth with an "mpls traffic-eng flooding thresholds" command, and identifying traffic engineering output fields that include administrative weight, reservable bandwidth, and bandwidth thresholds.); Ex.1024 (Entitled "Multiple Metrics for Traffic Engineering with IS-IS and OSPF" describing the metrics "administrative weight" and "bandwidth available."). Ex.1003, ¶155.

Doshi's disclosure is in the context of "multi-protocol label switched (MPLS) traffic engineering," which would have further confirmed a POSITA's understanding that Doshi's link utilization, utilization threshold, and administrative weight each correspond to traffic engineering metrics. Ex.1005, [0298], [0010], [0093], [0191], [0257]. Ex.1003, ¶¶152-157.

10. Claim 7

[7.1] ***The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).***

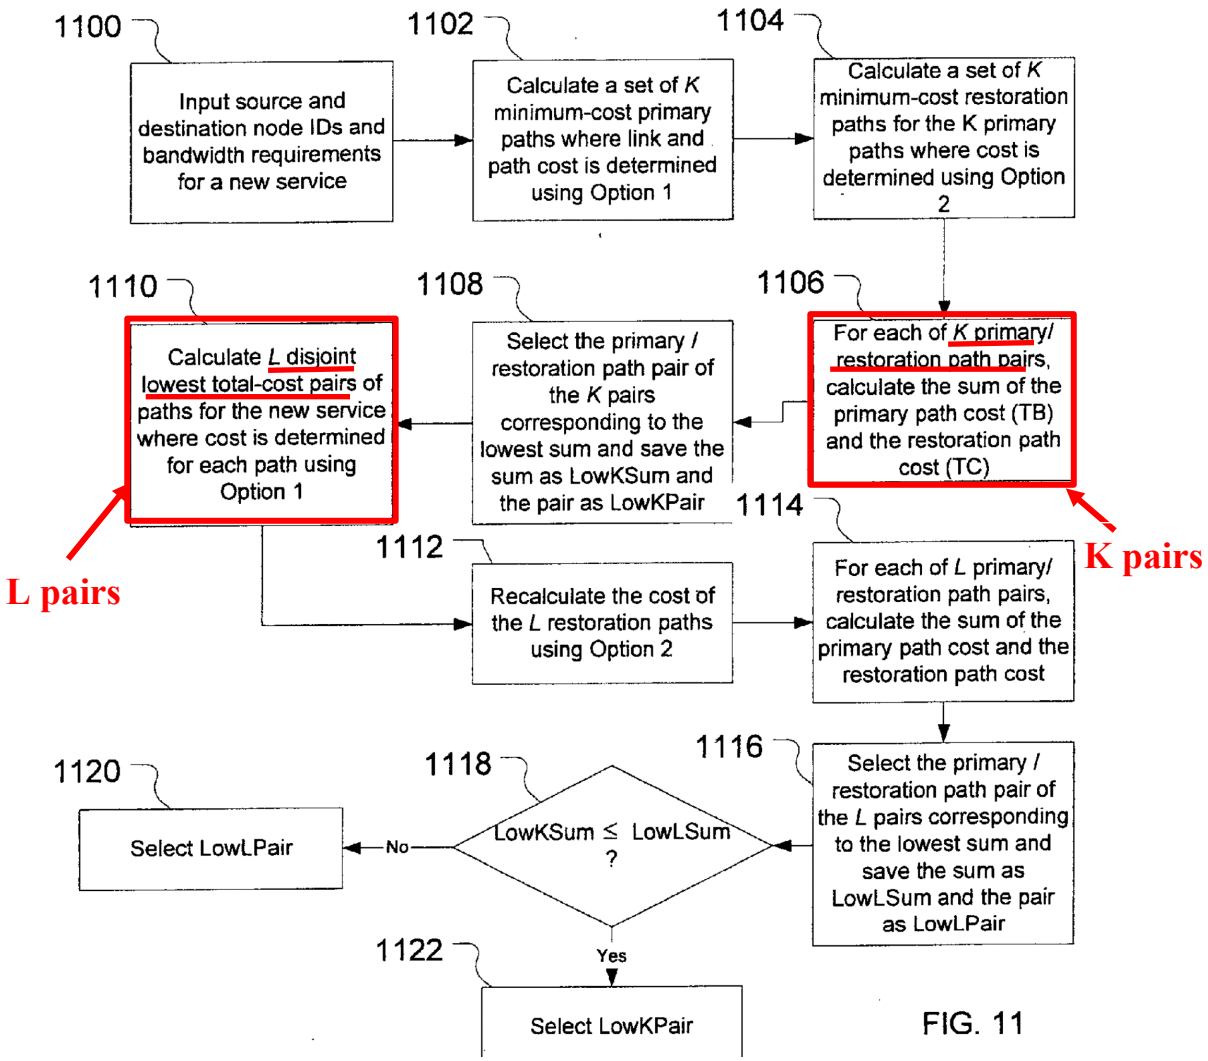
See claim 6. Ex.1003, ¶¶158-159.

11. Claim 9

[9.1] ***The method of claim 4, wherein said overall cost function comprises: selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and***

As discussed at [1.2]-[1.3], Doshi calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair and sums the path costs to find the cost of the pair together. The cost calculation considers various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness, and which represent the "*overall cost"* for each candidate path pair. The calculation used to determine the "*overall cost*" corresponds to the claimed "*overall cost function*." Ex.1003, ¶161.

As shown at Doshi's Figure 11 and disclosed in the accompanying description, the cost calculation includes calculating and comparing two sets of candidate path pairs—"K candidate pairs" and "**L candidate pairs**," where the L pairs are disjoint pairs. Ex.1005, [0142]-[0145]; Ex.1003, ¶162.



Ex.1005, Fig. 11 (annotated); Ex.1003, ¶162.

The L candidate pairs and K candidate pairs are both subsets of the total number of candidate pairs available. In the case of the L candidate pairs, the subset consists of pairs that are disjoint. As discussed at [5.1], the pairs include both a working path (e.g., working LSP) and a protection path (e.g., protection LSP) that are disjoint from each other, which substantially minimizes the probability of concurrent failure. Thus, Doshi's L candidate pairs correspond to "*a subset of*

entity pairs." A POSITA would have understood selecting the set of L candidate pairs for calculations corresponds to "*selecting*" as claimed. Ex.1003, ¶¶160-165.

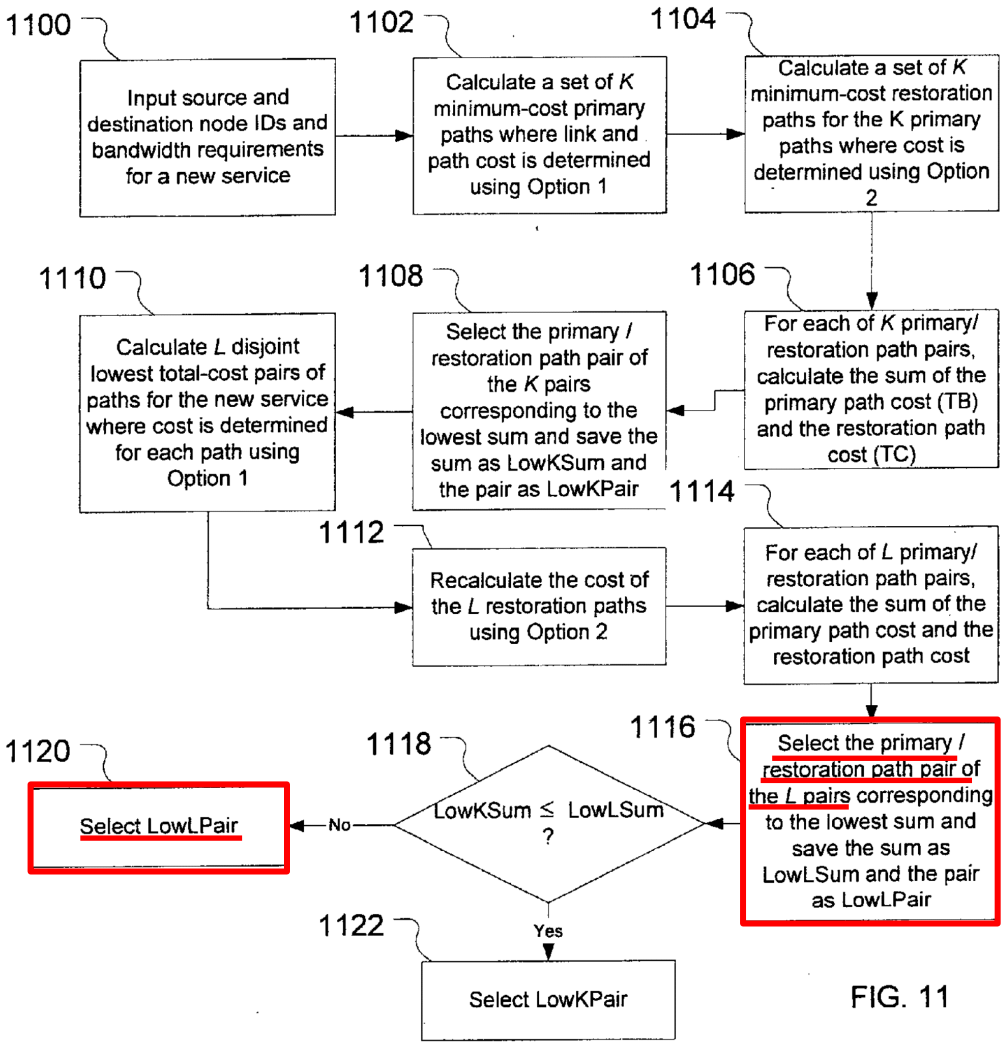


FIG. 11

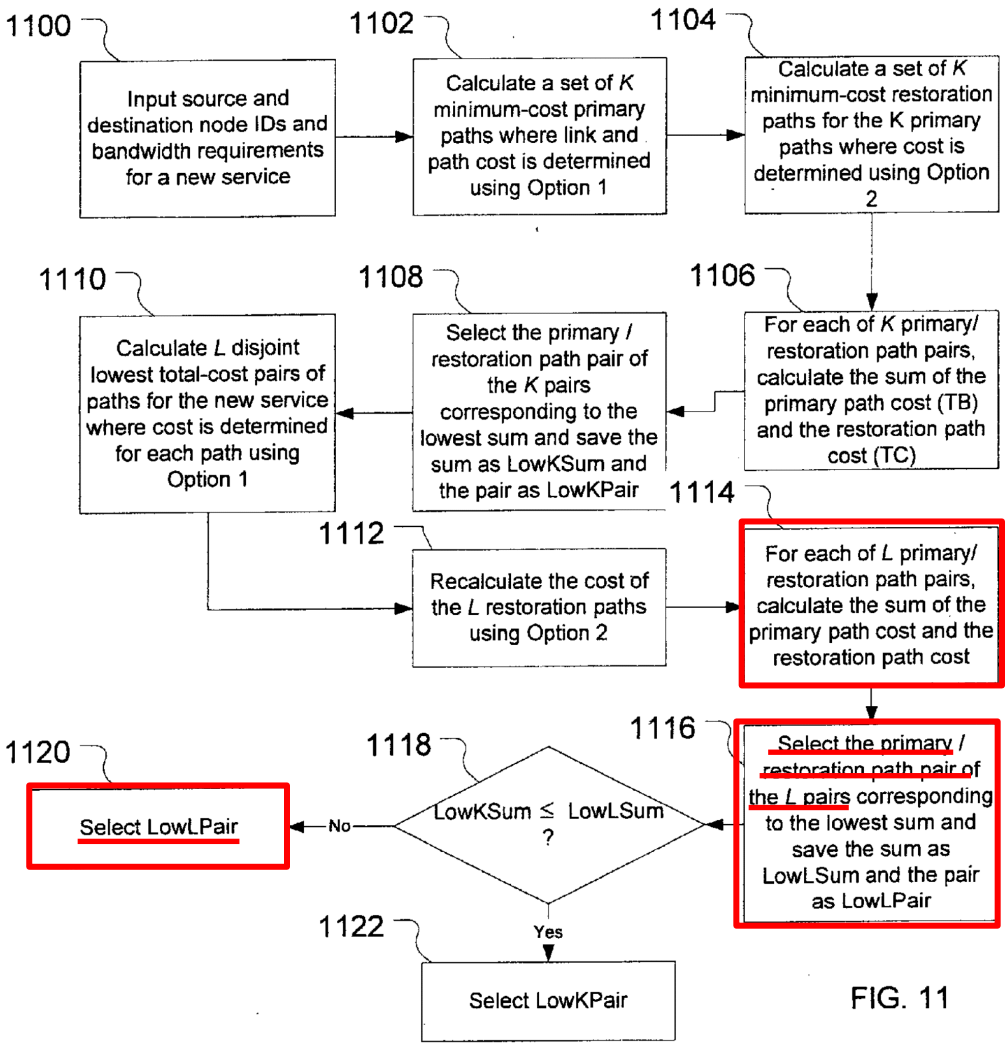
Ex.1005, Fig. 11 (annotated); Ex.1003, ¶164.

[9.2] ***if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.***

As discussed at [9.1], Doshi calculates L "disjoint lowest total-cost pairs," plural. Thus, Doshi's L candidate pairs include at least two pairs. Doshi goes on to

explain that the L candidate pairs are assessed for cost, which as discussed at [1.2] includes assessing the cost of each path in the pair. As discussed at [1.2]-[1.3] and [4.1], Doshi calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair. The cost calculation considers various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness. Ex.1005, [0124]-[0141]. The calculation used to determine the path cost corresponds to the claimed "*entity cost function*." Ex.1005, [0144]-[0145], Fig. 10 (element 1022 "Return PathCost"); Ex.1003, ¶166

Doshi describes selecting the lowest cost pair of the L candidate pairs as the lowest cost pair out of all pairs, including K candidate pairs. Ex.1005, [0142]-[0145]; Ex.1003, ¶167.



Ex.1005, Fig. 11 (annotated); Ex.1003, ¶167.

Accordingly, it would have been obvious to a POSITA that selection of the lowest cost L pair—the lowest cost pair out of all the pairs—"*minimizes an entity cost function*." Ex.1003, ¶¶ 166-169.

12. Claim 10

[10.1] ***The method of claim 9, wherein said entity cost function comprises a predefined metric.***

As discussed at [9.2], Doshi's cost calculation considers various factors,

including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness and corresponds to the claimed "*entity cost function*." Further, as discussed at [6.1], these factors, along with traffic engineering, correspond to a "*predefined entity cost metric*," which renders obvious "*a predefined metric."* Ex.1003, ¶¶170-171.

13. Claim 11

[11.1] ***The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).***

See claim 10. Ex.1003, ¶¶172-173.

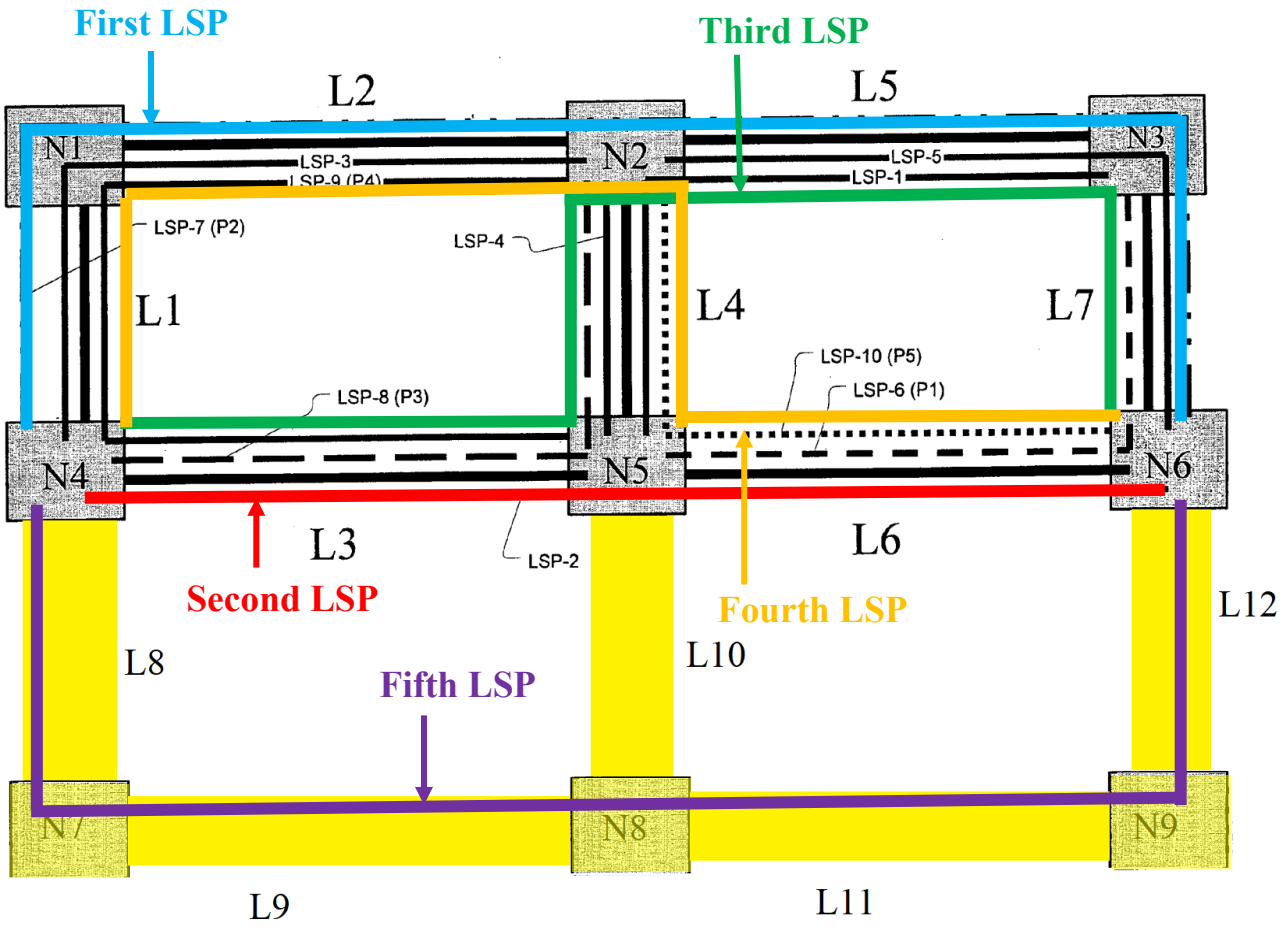
14. Claim 13

[13.1] ***The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity.***

Regarding replacement of the claimed "*working entity*" and "*protection entity*," as discussed at [2.1] and [5.1], Doshi's "primary path" is also referred to as a "working path" (e.g., working LSP), and its "restoration path" is also referred to as a "protection path" (e.g., protection LSP). Ex.1005, [0046]. Doshi's primary path (aka working path) corresponds to the claimed "*working entity*" and Doshi's restoration path (aka protection path) corresponds to the claimed and "*protection entity.*" Ex.1003, ¶175.

As discussed at [1.3]-[1.4], Doshi teaches that a primary and a protection pair (e.g., LSP pair) are selected from the plurality of paths of the MPLS network and that, in view of Guichard's teaching of ongoing reoptimization efforts, it would have been an obvious to perform Doshi's cost calculations repeatedly to attempt to determine whether the originally selected LSP pair still has the overall minimum-cost or if another LSP pair should be selected. Ex.1003, ¶176.

It would have been obvious for the LSPs of the originally selected pair (e.g., a primary LSP and a protection LSP) to be sequentially replaced as the ongoing reoptimization efforts identify lower cost LSPs. To build upon an earlier example, consider Doshi's modified and annotated Figure 1 below. Ex.1003, ¶177.



Ex.1005, Fig. 1 (modified and annotated); Ex.1003, ¶177.

As previously discussed at [1.3]-[1.4], the Fifth LSP was selected as the primary path and the Second LSP was selected as the protection path. Upon an event trigger as taught by Guichard, attempted reoptimization of the LSPs would take place to determine if there is a better or more optimal LSP. If, for example, the reoptimization evaluation determined that the links of the Fifth LSP are heavily utilized as compared to the links of the First LSP, it would have been obvious for Doshi's cost determination to result in a new pair (e.g., the First LSP as the

primary while the Second LSP remains as the protection) as having the overall minimum cost. Thus, the Fifth LSP would be replaced. A subsequent such reoptimization evaluation would find that, for example, the Second LSP has become suboptimal as a protection path and replace the Second LSP with another LSP e.g., a newly added Sixth LSP (not shown)). Of course, this is just an example, and given network status it would be understood that other LSPs may be selected as the overall minimum cost pair, but the LSPs have been replaced sequentially in this example—first the primary LSP followed by the protection LSP. It would have been obvious for the sequence of replacement to be reversed in other examples where the protection LSP becomes suboptimal before the primary LSP. Ex.1003, ¶178

It would have been similarly obvious for a single event-triggered reoptimization evaluation to find that both the Fifth and Second LSPs are suboptimal. When both LSPs are determined to be suboptimal, a POSITA would have found it obvious to replace the LSPs sequentially to minimize interruption in service. In that regard, a POSITA would have appreciated the desirability of maintaining a working and protection path throughout the replacement process. In replacing the Fifth and Second LSPs with the First and Sixth LSPs, for example, a POSITA would have found it obvious to first designate the First LSP (determined to be the optimal primary LSP) as a new protection LSP temporarily, thereby

replacing the Second LSP. Next, traffic on the Fifth LSP would be switched over to the First LSP with the Fifth LSP temporarily serving as a protection LSP. The First LSP would thus become the new primary LSP as was determined to be optimal. The Fifth LSP, now acting as a protection LSP, would then be replaced by the Sixth LSP, which was determined to be the optimal protection LSP. A POSITA would have recognized that such sequential replacement would advantageously maintain two LSPs (e.g., one primary and one protection) throughout the replacement. Ex.1003, ¶¶174-180.

15. Claim 17

Claim 17 is substantially similar to claim 1 and therefore is obvious for the reasons stated above. The below analysis addresses the obviousness of the preamble's different recitations, to the extent limiting. Ex.1003, ¶182.

[17.0] ***Non-transitory computer readable media configured to perform a method comprising the steps of:***

As discussed at [1.0], Doshi discloses that the claimed method steps are performed by a network manager, that may be a server or other computing device. It was common knowledge that computing devices, such as servers, would include a processor and memory. *See* Ex.1025 [0092], Fig.10 (describing "exemplary computer system 1000 that can be used as … a server" where the computer system 1000 includes "one or more processors," "a main memory," "and may also include

a secondary memory"); *see also* Ex.1031; Ex.1032; Ex.1033. Indeed, Doshi expressly discloses storing, for example, protection information and information about path pairs which would have informed a POSITA that Doshi's network manager includes a memory. Ex.1005, [0143]-[0144], [0196], [0214]; Ex.1003, ¶183.

It was also known for memory to be non-volatile and computer readable by the processor ("*[n]on-transitory computer readable media*"). *See* Ex.1016 (describing "computer-readable" hard disk drives and flash memory cards "provide nonvolatile storage of data, data structures, computer-executable instructions, and so forth."); *see also* Ex.1017; Ex.1018; Ex.1025, [0092]; *see also* Ex.1003, ¶184.

Accordingly, a POSITA would have found it obvious for Doshi's network manager to include a memory that corresponds to a "*non-transitory computer readable medium*." As discussed above at [1.0]-[1.5], Doshi's network manager performs the method steps. Since, as evidenced by the citations above, it was common knowledge to store executable instructions in memory, and since Doshi's network manager performs the method steps, a POSITA would have considered it obvious to store instructions to perform each step of the method in the memory of Doshi's network manager. Accordingly, a POSITA would have found it obvious for Doshi's network manager to include "*[n]on-transitory computer readable media configured to perform a method*." Ex.1003, ¶¶183-186.

[17.1]-[17.5]

See [1.1]-[1.5]. Ex.1003, ¶¶187-191.

16. Claim 18

[18.1] ***The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises: selecting a working entity from said plurality of transport entities; selecting a protection entity from said plurality of transport entities; and***

As discussed above at [1.3], Doshi discloses selecting an LSP pair comprising a primary LSP and a protection LSP from the plurality of MPLS LSPs. Doshi's "primary path" corresponds to a "working path." Ex.1005, [0046]. Thus, Doshi's step of selecting an LSP pair comprises selecting a "working path" and a protection path from the MPLS LSPs, which renders obvious this limitation. Ex.1003, ¶¶192-194.

[18.2] ***selecting an active entity from the set consisting of said working entity and said protection entity.***

See claim 3. Ex.1003, ¶195.

17. Claim 19

[19.1] ***The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.***

See claim 4. Ex.1003, ¶¶196-197.

18. Claim 20

[20.1] ***The non-transitory readable media of claim 19, wherein said overall cost function comprises: minimizing a probability of concurrent failure of said protection entity and said working entity; and***

See claim 5. Ex.1003, ¶¶198-199.

[20.2] ***a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).***

See claims 6 and 7. Ex.1003, ¶200.

C. Ground 2: Claims 8 and 12 are obvious under 35 U.S.C. § 103(a) over Doshi in view of Guichard and Huang.

1. Summary of Huang

Like the '821 patent and Doshi, Huang describes LSP selection within an MPLS network. Ex.1007, [0004], [0032], [0035], [0040]. Huang's LSP selection is described in the context of "protection of connections formed through a mesh-type communication network." Ex.1007, [0001]. In Huang, a request is received to set up an LSP segment between a head end node and a tail end node, where the request further identifies a backup route to the tail end node. Ex.1007, [0016]. The backup routes are LSPs that "protect each of the working links" if failure occurs. Ex.1007, [0032]. Huang's "working links" are also LSPs. *See* Ex.1007, [0030]-[0031]. Huang further describes "revertive switching," where "the connection will be switched back from the backup LSP to the working link once the working link has cleared the failure that caused the switchover." Ex.1007, [0053]; Ex.1003, ¶¶201-202.

2. Reasons to Combine Doshi and Huang

A POSITA would have been motivated to combine the teachings of Doshi and Huang, as discussed below. Ex.1003, ¶203.

First, Doshi and Huang are analogous art to the '821 patent. For example, just like the '821 patent, both references generally pertain to MPLS networks. Ex.1001, 3:15-18, Abstract; Ex.1005, [0049]; Ex.1007, [0004]-[0007]. Additionally, both references address the problem of selecting a pair of MPLS paths from available paths. Ex.1001, 4:63-5:8; Ex.1005, [0014]; Ex.1007, [0004], [0032], [0035], [0040], [0048]. Thus, given the similarities between the references and the fact that they address the problem of selecting LSPs, a POSITA considering Doshi would have naturally considered the teachings of Huang. Ex.1003, ¶204.

Second, a POSITA would have been motivated to combine the teachings of Doshi and Huang to produce numerous predictable and beneficial results. Ex.1003, ¶205.

Doshi teaches "Switching Between Working and Protection LSPs" and explains that "after detecting a failure, end nodes of an LSP switch traffic from a primary (i.e., working) LSP to its corresponding protection LSP." Ex.1005, [0266]-[0267]. Huang complements Doshi by teaching what to do when the failure affecting the working LSP has been cleared. Huang teaches "revertive switching"

where "the connection will be switched back from the backup LSP to the working link once the working link has cleared the failure that caused the switchover." Ex.1007, [0053]; Ex.1003, ¶206.

A POSITA would have been motivated to configure Doshi's working LSP as revertive, in view of Huang, to avoid the need to calculate and select a replacement working LSP, when the previously selected working LSP has been quickly restored. Reverting allows for the previously selected working LSP (which as discussed at [1.2]-[1.4] is part of the overall minimum-cost pair) to be utilized without any additional computations or selections. A POSITA would have recognized that in cases of quick restoration, e.g., when a mistaken failure diagnosis is quickly resolved, it would be simpler to revert back to the working LSP instead of having the network manager calculate the costs for all LSPs to ultimately reach the same result of selecting the working LSP. Thus, Huang's teachings obviate unnecessary efforts. Ex.1003, ¶¶207-208.

The combination is merely the use of a known technique (Huang's technique of configuring the working LSP as revertive) to improve a similar method (Doshi's method for selecting working and protection LSPs) in the same way with predictable results (reverting back to Doshi's working LSP once the failure has been cleared without performing needless calculations). Ex.1003, ¶209.

The results would have been predictable and there would have been a reasonable expectation of success in the combination. A POSITA would have had a reasonable expectation of success because Huang provides implementation details, such as waiting to revert back until the working LSP has been restored. Ex.1007, [0053]. Also, the expectation of success is evidenced by the fact that revertive switching was in described in multiple RFCs discussing MPLS operation. Ex.1011, 12; Ex.1030, 16; Ex.1003, ¶¶210-212.

Implementing the combination would have been within a POSITA's skillset since Doshi already sets forth techniques for switching between working and protection LSPs. Ex.1005, [0266]-[0268]. Building upon Doshi to configure working LSPs as revertive, per Huang, would utilize already existing switching techniques. Doshi acknowledges that a skilled artisan would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Ex.1005, [300]. It would have been obvious for such modifications to include reverting back to the working LSP after a failure of the working LSP has been cleared. Ex.1003, ¶¶211-213.

3. Claim 8

[8.1] ***The method of claim 4, further comprising the step of configuring said working entity as revertive.***

First, as discussed at [2.1], Doshi describes "Switching Between Working

and Protection LSPs" and explains that "after detecting a failure, end nodes of an LSP switch traffic from a primary (i.e., working) LSP to its corresponding protection LSP." Ex.1005, [0266]-[0267]. Doshi also incorporates by reference in its entirety RFC 3209 (Ex.1005, [0056]), which describes returning to an "original path when the failed resource becomes re-activated." Ex.1011, 12. Since Doshi describes switching from a working LSP to a protection LSP after a failure, and since Doshi incorporates RFC 3209's disclosure of returning to an original path that is re-activated after a failure, it would have been obvious to a POSITA to configure the working LSPs as revertive. Thus, Doshi renders obvious "*configuring said working entity as revertive*." Ex.1003, ¶¶214-215.

Second, like RFC 3209, Huang also teaches "that, in the event of a failure" of a working link, traffic "may be switched to corresponding individual backup LSPs" (e.g., a "protection LSP" in terms of Doshi). Ex.1007, [0032]. The "working links" are LSPs. *See* Ex.1007, [0030]-[0031]; Ex.1003, ¶¶216-217. Huang also discloses "**revertive switching,**" such that the "the connection will be switched back from the **backup LSP to the working link** once the working link has cleared the failure that caused the switchover," which corresponds to "*configuring said working entity as revertive.*" Ex.1007, [0053]; Ex.1003, ¶218.

A POSITA would have found it obvious to utilize revertive switching in Doshi, as taught by Huang, such that the switch configures the working LSP to be

used again once the fault is cleared. A POSITA would have recognized that in some instances, e.g., in the case of a quickly resolved mistaken failure diagnosis, it would be simpler to revert back to the working LSP instead of calculating the costs for all LSPs only to again select the working LSP. *See also* Reasons to Combine Doshi and Huang, § VIII.C.2. Ex.1003, ¶¶214-220.

4. Claim 12

[12.1] ***The method of claim 10, further comprising the step of configuring said working entity as revertive.***

See claims 8 and 10. Ex.1003, ¶¶221-222.

D. Ground 3: Claims 14-16 are obvious under 35 U.S.C. § 103(a) over Doshi in view of Guichard and Xu.

1. Summary of Xu

Like the '821 patent and Doshi, Xu describes selecting paths to protect against failures in an MPLS network. Ex.1025, [0001], [0005]-[0006], [0030], [0035]-[0038], [0043]. Xu discloses "selecting the end-to-end paths" based on a "set of modeled failure states." Ex.1025, [0006], [0043], [0049]. The states represent "link failures that occur simultaneously" with "a predetermined probability of occurrence." Ex.1025, [0006]-[0007], [0049], [0053], [0080]; Ex.1003, ¶¶223-225.

2. Reasons to Combine Doshi and Xu

A POSITA would have been motivated to combine the teachings of Doshi and Xu, as discussed below. Ex.1003, ¶226

First, Doshi and Xu are analogous art to the '821 patent. Just like the '821 patent and Doshi, Xu generally describes MPLS networks. Ex.1001, 3:15-18, Abstract; Ex.1005, [0049]; Ex.1025, [0036], [0038]. Xu also addresses the problem of selecting paths within an MPLS network. Ex.1001, 4:63-5:8; Ex.1005, [0014]; Ex.1025, [0002], [0030], [0040], [0043], [0060]. Given the similarities between the references and their addressing the same problem of selecting paths, a POSITA considering Doshi would have naturally considered the teachings of Xu. Ex.1003, ¶227.

Second, a POSITA would have been motivated to combine the teachings of Doshi and Xu to produce numerous predictable and beneficial results. Ex.1003, ¶228.

Doshi teaches protecting services against link failures by considering disjointedness when selecting primary and protection paths. Ex.1005, [0155]-[0156]. Strictly disjoint paths share no common links or nodes other than their common ingress and egress nodes, so the failure of a single network element would not affect both paths. Ex.1005, [0055], [0158]. Doshi recognizes, however, that there "will be some cases … where no strictly disjoint path exists" or where

"other factors…may lead to selection of a protection path, for a given primary path, that is not strictly disjoint from the primary path." Ex.1025, [0055]. It would be obvious, in such instances, to at least minimize the likelihood of a concurrent failure. Xu complements Doshi by teaching that a probability of occurrence is determined for each of a set of failure states and that path selection decisions consider the probability of various failure states occurring. Ex.1025, [0006], [0053]; Ex.1003, ¶229.

A POSITA would have been motivated to determine a probability of occurrence of failure states, as taught by Xu, when implementing Doshi's path selection in order to further inform selection decisions and better protect services against link failures in circumstances where no strictly disjoint pair of paths exists or where other factors lead to selection of a partially disjoint pair as discussed above. Determining the probability of various failure states would further refine selection between pairs of paths that have the same disjointedness (e.g., between pairs all having a primary path that shares one link in common with its protection path). Such pairs have the same disjointedness, but there may be differences in risk of failure for different links. A pair in which the link in common between primary and protection paths of the pair has a higher risk of failure would, everything else being equal, be more likely to concurrently fail than a pair in which the link in common has a lower risk of failure. A POSITA would have been motivated to

apply Xu's probability determination to differentiate between such pairs to improve reliability by reducing the probability of concurrent failure of Doshi's selected pair. Ex.1003, ¶230.

The combination is merely the use of a known technique (Xu's probability determination) to improve a similar method (Doshi's method for selecting paths) in the same way with predictable results—selecting paths with lower probability of failure. The results would have been predictable, and there would have been a reasonable expectation of success in the combination since Doshi and Xu are analogous art in the same field of endeavor. Ex.1003, ¶231.

Implementing the combination would have been within a POSITA's skillset because techniques for determining the probability of network failures were known in the art. *See* Ex.1026. Xu leaves implementation details of failure probability determination to a POSITA, recognizing that a POSITA would have had sufficient skills to configure the determination. Xu explains that "the present disclosure [is] not limited to any specific combination of hardware and software and the computer program code required to implement the foregoing can be developed by a person of ordinary skill in the art." Ex.1025, [0094]-[0096]. Xu's techniques can be "practiced or carried out in a variety of applications." Ex.1025, [0029]. Since Doshi is analogous art in the same field of endeavor, it would have been obvious to carry out Xu's probability determination in the context of Doshi's path selection

with a reasonable expectation of success. Indeed, Doshi acknowledges that a POSITA would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Ex.1005, [300]. It would have been obvious for such modifications to include Xu's failure state probability determination. Ex.1003, ¶¶232-233.

3. Claim 14

[14.0] ***A system for selecting entities within an MPLS network, comprising:***

As discussed at [1.0], Doshi discloses selecting LSPs within an MPLS network, which renders obvious "*selecting entities within an MPLS network*." Doshi's network components, e.g., network manager, nodes, modules, switches, etc., and corresponding communication paths, that are involved in the selecting correspond to the claimed "*system*." *See* Ex.1005, [0014], [0035], [0081], [0214], [0271]; Ex.1003, ¶¶234-235.

[14.1] ***a data structure comprising a plurality of transport entity descriptors;***

First, as discussed at [1.1], Doshi in combination with Guichard discloses providing a plurality of LSPs between two end nodes where the LSPs correspond to "*a plurality of [] transport entit[ies].*" Ex.1003, ¶236.

Second, as discussed at [1.2]-[1.5], Doshi in combination with Guichard discloses calculating the cost of various LSP pairs and selecting the lowest cost pair. A POSITA would have found it obvious to store and reference "*descriptors*"

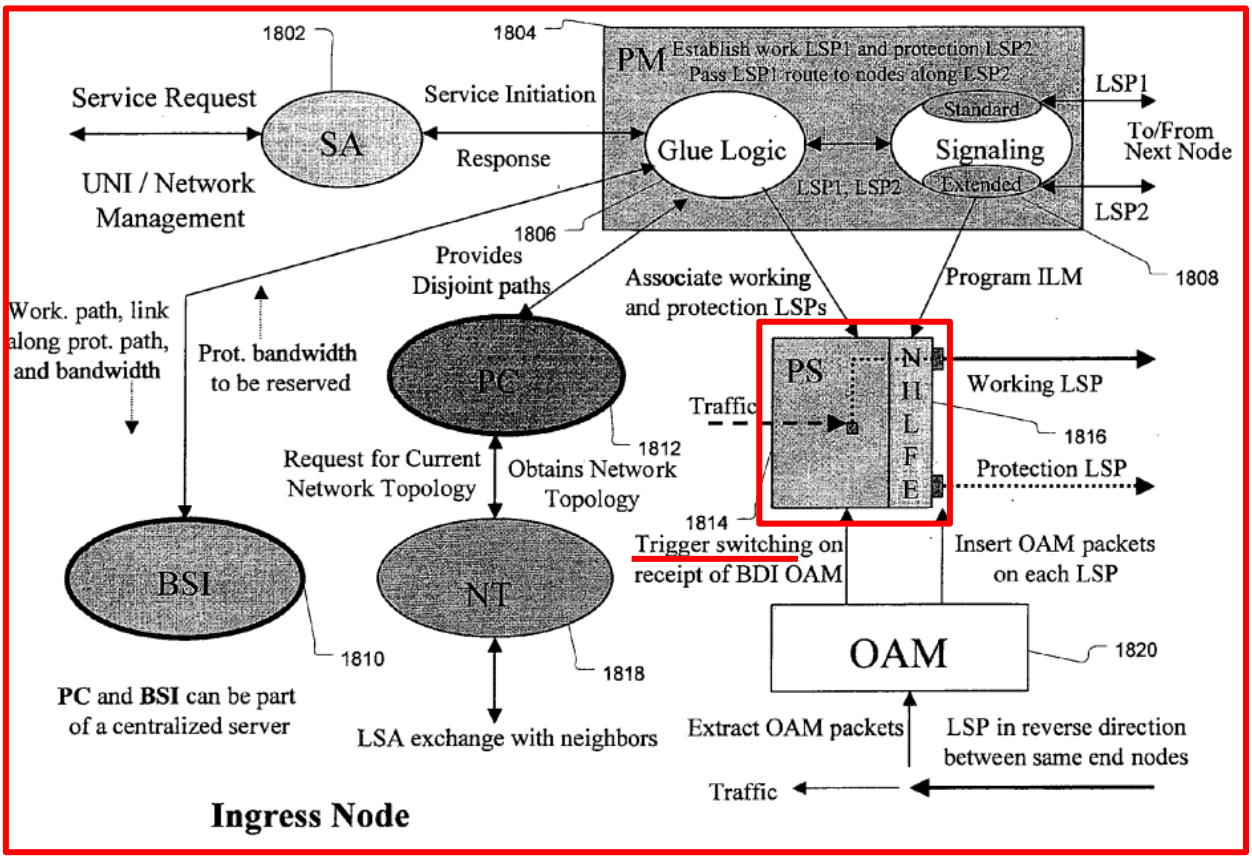
of the LSPs at least temporarily to accomplish the cost calculation and selection described above. Such storage would be in a memory, entries in which a POSITA would understand corresponds to a "*data structure*." Ex.1003, ¶237

Moreover, Doshi expressly describes saving information about a plurality of LSP pairs in a data structure during cost calculation. Ex.1005, [0143] ("[T]he minimum-cost pair is saved in the **data structure** LowKPair."); Ex.1005, [0144] ("[T]he lowest-cost pair is stored in the **data structure** LowLPair."). Doshi also maintains "an extended link-state database" (referred to as "TE/Share database") that include information used for "path computation." Ex.1005, [0196]. "This database contains a **data structure** for each link L that the node owns. The data structure for each link contains information (e.g., bandwidth, link-id) about all the other links in the network for which link L provides restoration capacity." Ex.1005, [0197]; *see also* Ex.1005, [0082], [0091]-[0095], [0172]-[0176]; Ex.1003, ¶¶238-239.

Doshi also describes a functional architecture that includes a "Path Management" module that is "responsible for path setup, refresh, tear-down, and monitoring functions." Ex.1005, [0273]. It would have been obvious to a POSITA to store "*descriptors*" of the monitored paths in a "*data structure*" to differentiate the paths and to provide data entries for the paths such that changes can be tracked over time and saved in association with the relevant path. Ex.1003, ¶¶240-241.

[14.2] ***an entity protection switch configured to switch between a working entity and a protection entity; and***

Doshi describes that "after detecting a failure, **end nodes of an LSP switch traffic from a primary (i.e., working) LSP to its corresponding protection LSP**." Ex.1005, [0266]-[0267]; *see also* [0266] ("Switching Between Working and Protection LSPs"). Doshi further discloses a "**Protection Switching (PS)**" module that is "**responsible for switching the affected traffic onto a protection LSP** after detecting a failure or receiving failure notification." Ex.1005, [0278]. Doshi's exemplary Protection Switching module at Figure 18 illustrates its inclusion in a node of an LSP. Ex.1003, ¶242.



Ex.1005, Fig. 18 (annotated); Ex.1003, ¶242.

Thus, Doshi's node (including the Protection Switching module) or the Protection Switching module itself, which is configured to switch between the working LSP and the protection LSP, renders obvious this limitation. Ex.1003, ¶¶242-244.

[14.3] ***digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:***

First, as discussed at [1.0]-[1.5] and [2.1], Doshi in combination with Guichard discloses logical flowcharts (e.g., at Figures 10 and 11) for selecting a working primary LSP ("*working entity*") and a protection LSP ("*protection*

entity"). As discussed at [14.1], a POSITA would have found it obvious to store and reference "*descriptors*" of the LSPs to accomplish cost calculation and selection of the LSPs. Accordingly, selection of a working LSP and a protection LSP from the descriptors renders obvious "*select[ing],*" as claimed. Ex.1003, ¶245.

Second, a POSITA would have found it obvious to implement Doshi's flowchart logic (e.g., at Figures 10 and 11) with software or a combination of software and hardware ("*digital logic*") to perform the calculation and selection techniques. Doshi discloses a series of modules for performing its calculation and selection techniques. Ex.1005, [0270]-[0281]. A POSITA would have found it obvious to implement these modules in, for example, software or a combination of software and hardware ("*digital logic*") since it was well known that such an implementation would allow for execution of instructions to control functionality (e.g., performing calculations and making decisions). *See* Ex.1016 (describing "computer-executable instructions for performing the methods of the disclosed innovation."); Ex.1017 (describing "software and/or firmware provided in a read only memory … for execution" by a processor "to implement the various functions as detailed below."); Ex.1018 (describing using "software instructions to implement the present invention."). Software or a combination of software and hardware that executes instructions, for example to carry out the flowchart of

Figures 10 and 11, to calculate and select the lowest cost pair corresponds to "*digital logic,*" as claimed. Ex.1003, ¶246.

A POSITA would have found it obvious to include such software or a combination of software and hardware, for example, in Doshi's "network manager," including the "Path Management" module that is "responsible for path setup, refresh, tear-down, and monitoring functions." *See* Ex.1005, [0271]-[0273]; Ex.1003, ¶¶247-248.

[14.4] ***logic configured to determine a probability of concurrent failure of said working entity and said protection entity;***

First, as discussed at [1.2], [2.1], [5.1], Doshi's path calculation considers disjointedness of the working LSP ("*working entity*") and the protection LSP ("*protection entity*"). Further, as discussed at [14.3], a POSITA would have found it obvious to use "*digital logic*" to perform Doshi's flowchart logic, including path calculation and selection techniques. Ex.1003, ¶¶249-250.

Doshi recognizes that when two paths are disjoint, e.g., by being node- and link-disjoint, "then a failure affecting one of them will not affect the other." Ex.1005, [0040]. A POSITA would have understood that probability of concurrent failure is inversely related to disjointedness—as paths become more disjoint, the likelihood they will fail concurrently goes down. Doshi recognizes that when two paths are disjoint, e.g., by being node- and link-disjoint, "then a failure affecting

one of them will not affect the other." Ex.1005, [0040]. It would take two distinct failures (one affecting each path) for both paths to fail concurrently. Since Doshi explains that the "probability of occurrence" of two such failures is "**very insignificant**," a POSITA would have understood that determining that two paths are disjoint with no nodes or links in common is also a determination that the paths have a "very insignificant" probability of failing concurrently. Ex.1005, [0040]. Accordingly, determining path disjointedness (as part of Doshi's flowchart) corresponds to "*determine[ing] a probability of concurrent failure*" since knowledge of, e.g., strict disjointedness represents a determination that there is a "very insignificant" probability that the paths will fail at the same time. Ex.1003, ¶¶250-251.

Second, Xu discloses a "network architecture [that] uses multiple paths between each ingress-egress router pair" in an MPLS network and provides techniques for "selecting the end-to-end paths" between the pair by "Network-Management Software (NMS)." Ex.1025, [0030]-[0038]. To inform path selection, Xu models failure states that would cause at least one path to fail. Ex.1025, [0006], [0049], [0080], Claim 1, Claim 11. A POSITA would have understood that at least some of the modeled failure states would be states in which the failure affects two paths simultaneously—corresponding to a state of concurrent failure. Ex.1003, ¶¶252-253.

Xu explains that each failure state has a "predetermined probability of occurrence." Ex.1025, [0006]. A POSITA would have understood that a "pre**determined**" probability has been "*determine[d].*" When a failure state is one where the failure affects two paths simultaneously—a state of concurrent failure—Xu's predetermining a probability of occurrence for that failure state corresponds to "*determin[ing] a probability of concurrent failure*." Each failure state is assigned a weight in Xu's selection function based on how common the state is (probability of occurrence). Ex.1025, [0010]; *see also* [0033], [0045], [0053], [0080]. Ex.1003, ¶254.

As discussed in more detail in VIII.D.2, a POSITA would have found it obvious determine the probability of occurrence of failure states, as taught by Xu, when implementing Doshi's calculation and selection techniques and would have been motivated to determine the probability of various failure states to achieve a more refined selection between pairs of paths that have the same disjointedness. Ex.1003, ¶¶255-257.

[14.5] ***logic configured to determine an entity cost of said plurality of transport entity descriptors: and***

First, as discussed at [1.2], Doshi's calculates the cost for each LSP path of candidate LSP pairs. This is done for all candidate LSP pairs, which means that a cost is calculated for each LSP path. As discussed at [14.1], a POSITA would have

found it obvious to store and reference "*descriptors*" of the LSPs at least temporarily to accomplish cost calculation and selection of the LSPs. Further, as discussed at [14.3], a POSITA would have found it obvious to use "*digital logic*" to perform Doshi's flowchart logic, including path calculation and selection techniques. Ex.1003, ¶¶258-260.

[14.6] ***logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,***

First, as discussed at [1.0]-[1.5] and [2.1], Doshi in combination with Guichard discloses reselecting an LSP pair from the plurality of MPLS LSPs in response to an event. The reselection includes reselecting a primary LSP and a protection LSP if they remain the lowest cost pair. Doshi's "primary" LSP is also referred to as a "working path" (e.g., working LSP). Put differently, Doshi discloses reselecting a working LSP and a protection LSP. As discussed above at [2.1], these working and protection LSPs correspond to the claimed "*working entity*" and "*protection entity,*" respectively. As discussed at [14.1], a POSITA would have found it obvious to store and reference "*descriptors*" of the LSPs at least temporarily to accomplish cost calculation and selection of the LSPs. Further, as discussed at [14.3], a POSITA would have found it obvious to use "*digital logic*" to perform Doshi's flowchart logic, including path calculation and selection techniques. Ex.1003, ¶¶261-263.

[14.7] ***wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.***

See [1.5]. Ex.1003, ¶264.

4. Claim 15

[15.1] ***The system of claim 14, wherein said entity protection switch comprises a 1:1 switch.***

As discussed at [14.2], Doshi discloses the claimed "*entity protection switch*." Ex.1003, ¶¶265-267. Doshi further discloses that "[i]f the bandwidth were allocated in advance of a failure, this would correspond, in the parlance of the field of protection and restoration for optical transport networks, to **a 1:1 protection scheme**." Ex.1005, [0060]. Doshi's protection switching module in the context of a 1:1 protection scheme renders obvious this limitation. Ex.1003, ¶¶265-268.

5. Claim 16

[16.1] ***The system of claim 14, wherein said entity protection switch comprises a 1+1 switch.***

As discussed at [14.2], Doshi discloses the claimed "*entity protection switch*." Ex.1003, ¶271. Doshi further discloses that "[i]f the bandwidth were not only allocated, but additionally if a copy of the service path's data were to be duplicated to the protection path, this would correspond to a **1+1 protection scheme**." Ex.1005, [0060]. In some circumstances, "**a 1+1 restoration scheme is**

the best option available for restoration" and therefore are "popular." Ex.1005, [0147], [0178]. Doshi's protection switching module in the context of a 1+1 protection scheme renders obvious this limitation. Ex.1003, ¶¶269-272.

IX. DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE

A. Discretionary denial under the *Fintiv* factors is not appropriate

The six factors considered for § 314 denial strongly favor institution. *See Apple Inc. v. Fintiv, Inc*., IPR2020-00019, Paper 11 (PTAB Mar. 20, 2020) (precedential).

1. No evidence regarding a stay

No motion to stay has been filed, so the Board should not infer the outcome of such a motion. *Sand Revolution II LLC v. Continental Intermodal Group – Trucking LLC*, IPR2019-01393, Paper 24 at 7 (PTAB June 16, 2020) (informative); *see also Dish Network L.L.C. v. Broadband iTV, Inc.*, IPR2020-01359, Paper 15 (Feb. 12, 2021). Thus, this factor is neutral.

2. Parallel proceeding trial date

This factor weighs strongly against discretionary denial because the projected trial date—based on median time-to-trial statistics—is in August of 2024,

after the Board's Final Written Decision is expected in July of 2024.[3] While trial is currently scheduled for March 4, 2024 (Ex.1014), the Board recognizes "that scheduled trial dates are unreliable and often change." *See* Director's June 21, 2022 Memorandum on Discretionary Denials ("Memo"), 8. "The PTAB will weigh this factor [factor 2] against exercising discretion to deny institution under Fintiv if the median time-to-trial is around the same time or after the projected statutory deadline for the PTAB's final written decision." Memo, 9.

The co-pending district court case was filed in the Eastern District of Texas on July 22, 2022. *See* Ex.1012. The most recent statistics show a median time-to-trial in the Eastern District of Texas at 24.5 months. Ex.1013, 5. Accordingly, the projected trial date for purposes of *Fintiv* is August of 2024—approximately 24 months after July 2022, and after the Board's Final Written Decision is expected in July of 2024. Because the projected trial date is "around the same time or after" the Board's expected final written decision, this factor weighs in favor of institution.

3. Investment in the parallel proceeding

The co-pending litigation is in its very early stages, and the investment in it has been minimal. The parties have not exchanged preliminary positions on claim

[3] July 2024 is 18 months after January 2023, when Petitioner expects a notice of accorded filing date for this petition.

construction or invalidity, expert discovery has not begun, and the parties have not exchanged their first set of discovery requests. *See PEAG LLC v. Varta Microbattery GmbH*, IPR2020-01214, Paper 8, 17 (Jan. 6, 2021). Further, the Markman hearing is not scheduled until September of 2023, two months after an expected institution decision by the Board. Ex.1014, 3.

Moreover, Petitioner only learned which claims were being asserted on November 3, 2022. *See* Ex.1015. Under *Fintiv*, Petitioner's prompt filing "weigh[s] against exercising the authority to deny institution." *Fintiv*, Paper 11 at 11 ("If the evidence shows that the petitioner filed the petition expeditiously, such as promptly after becoming aware of the claims being asserted, this fact has weighed against exercising the authority to deny institution under NHK"). This factor favors institution.

4. Overlapping issues with the parallel proceeding

There is no present overlap of prior art issues due to the early stage of district court litigation. For example, Petitioner has not served its preliminary invalidity contentions in the district court proceeding. Consequently, this factor favors institution.

5. Identity of parties

Petitioner is a defendant in the litigation. That is true of most Petitioners in IPR proceedings. Accordingly, this factor should not be a basis for denying institution.

6. Other circumstances

As discussed in detail above, the prior art presented in this Petition renders the Challenged Claims unpatentable as obvious. The merits of Petitioner's arguments are strong, and this factor weighs against discretionary denial. Memo, 4.

As such, because the *Fintiv* factors are either neutral or weigh against discretionary denial, and institution should not be denied on discretionary factors.

B. Discretionary denial under 35 U.S.C. § 325(d) is not appropriate

None of the references presented in the petition were cited or considered by the Examiner during prosecution of the '821 patent. Accordingly, discretionary denial under 35 U.S.C. § 325(d) is not appropriate.

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

The '821 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 (PTAB Sept. 6, 2016) (Section II.B.4.i. precedential).

X. CONCLUSION

Accordingly, Petitioner has established a reasonable likelihood that the Challenged Claims are unpatentable.

Respectfully submitted,

Dated: January 9, 2023

/Theodore M. Foster/

HAYNES AND BOONE, LLP
2323 Victory Avenue, Suite 700
Dallas, Texas 75219
Customer No. 27683

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

XI. 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 '821 patent is or was involved in the following case:

Case Heading	Number	Court	Date
Orckit Corporation v. Cisco Systems, Inc.	2-22-cv-00276	EDTX	Jul. 7, 2022

C. Lead and Back-up Counsel and Service Information

Lead Counsel
Theodore M. Foster
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219
Phone: (303) 382-6205
Fax: (214) 200-0853
ipr.theo.foster@haynesboone.com
USPTO Reg. No. 57,456

Back-up Counsel
David L. McCombs
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219
Phone: (214) 651-5533
Fax: (214) 200-0853
david.mccombs.ipr@haynesboone.com
USPTO Reg. No. 32,271

Gregory P. Huh
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219
Phone: (972) 739-6939
Fax: (214) 200-0853
gregory.huh.ipr@haynesboone.com
USPTO Reg. No. 70,480

Calmann J. Clements
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219

Phone: (972) 739-8638
Fax: (214) 200-0853
calmann.clements.ipr@haynesboone.com
USPTO Reg. No. 66,910

Please address all correspondence to lead and back-up counsel. Petitioner consents to service in this proceeding by email at the addresses above.

XII. CLAIMS APPENDIX

[1.0] 1. An entity selection method performed by a network device, comprising the steps of:

[1.1] providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;

[1.2] determining an overall cost for each entity pair of said plurality of entities;

[1.3] selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and

[1.4] if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,

[1.5] wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

[2.1] 2. The method of claim 1, wherein said step of selecting an entity pair further comprises: selecting a working entity from said entity pair; and selecting a protection entity from said entity pair.

[3.1] 3. The method of claim 2, further comprising the step of selecting an active

entity from the set consisting of said working entity and said protection entity.

[4.1] 4. The method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.

[5.1] 5. The method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.

[6.1] 6. The method of claim 4, wherein said overall cost function comprises a predefined entity cost metric.

[7.1] 7. The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).

[8.1] 8. The method of claim 4, further comprising the step of configuring said working entity as revertive.

[9.1] 9. The method of claim 4, wherein said overall cost function comprises:

selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and

[9.2] if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.

[10.1] 10. The method of claim 9, wherein said entity cost function comprises a predefined metric.

[11.1] 11. The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).

[12.1] 12. The method of claim 10, further comprising the step of configuring said working entity as revertive.

[13.1] 13. The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity.

[14.0] 14. A system for selecting entities within an MPLS network, comprising:

[14.1] a data structure comprising a plurality of transport entity descriptors;

[14.2] an entity protection switch configured to switch between a working entity and a protection entity; and

[14.3] digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:

[14.4] logic configured to determine a probability of concurrent failure of said working entity and said protection entity;

[14.5] logic configured to determine an entity cost of said plurality of transport entity descriptors: and

[14.6] logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,

[14.7] wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

[15.1] 15. The system of claim 14, wherein said entity protection switch comprises a 1:1 switch.

[16.1] 16. The system of claim 14, wherein said entity protection switch comprises a 1+1 switch.

[17.0] 17. Non-transitory computer readable media configured to perform a method comprising the steps of:

[17.1] providing a plurality of MPLS transport entities between a first endpoint and a second endpoint;

[17.2] determining an overall cost for each entity pair of said plurality of entities;

[17.3] selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and

[17.4] if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising

at least one entity distinct from the entities of said entity pair,

[17.5] wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

[18.1] 18. The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises: selecting a working entity from said plurality of transport entities; and selecting a protection entity from said plurality of transport entities; and

[18.2] selecting an active entity from the set consisting of said working entity and said protection entity.

[19.1] 19. The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.

[20.1] 20. The non-transitory readable media of claim 19, wherein said overall cost function comprises: minimizing a probability of concurrent failure of said protection entity and said working entity; and

[20.2] a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioner hereby certifies, in accordance with and in 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 13,961. 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 9, 2023

/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 9, 2023
Manner of service	PRIORITY EXPRESS MAIL
Documents served	Petition for *Inter Partes* Review Under 35 U.S.C. § 312 and 37 C.F.R. § 42.104 of U.S. 8,830,821; Petitioner's Exhibit List; All Exhibits; Petitioner's Power of Attorney.
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

EXHIBIT 3

Filed on Behalf of: Cisco Systems, Inc.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,

Petitioner,

- vs. -

ORCKIT IP, LLC,

Patent Owner

PETITION FOR INTER PARTES REVIEW

OF U.S. PATENT NO. 10,652,111

Case No.: IPR2023-00554

TABLE OF CONTENTS

I. INTRODUCTION 1

II. GROUNDS FOR STANDING 1

III. REQUESTED RELIEF AND REASONS FOR REQUESTED RELIEF 2

IV. OVERVIEW OF THE '111 PATENT 2

A. Summary of the '111 Patent 2

B. Priority Date 4

V. STATUTORY GROUNDS FOR CHALLENGES 4

VI. LEVEL OF ORDINARY SKILL IN THE ART 5

VII. CLAIM CONSTRUCTION 5

VIII. THE ART PRIOR TO THE '111 PATENT 8

A. Lin 8

B. Shieh 10

C. Swenson 12

IX. GROUND 1: CLAIMS 1-9, 12-24 AND 27-31 ARE UNPATENTABLE AS OBVIOUS OVER LIN IN VIEW OF SWENSON. 13

A. Claim 1 13

B. Claim 2 33

C. Claim 3 34

D. Claim 4 36

E. Claim 5 37

F. Claim 6 39

G. Claim 7 40

H. Claim 8 41

I. Claim 9 41

J. Claim 12 43

K. Claim 13 43

L. Claim 14 43

M. Claim 15 44

N. Claim 16 44

O. Claim 17 45

P. Claim 18 46

Q. Claim 19 47

R. Claim 20 47

S. Claim 21 49

T. Claim 22 49

U. Claim 23 50

V. Claim 24 50

W. Claim 27 52

X. Claim 28 52

Y. Claim 29 52

Z. Claim 30 53

AA. Claim 31 54

X. GROUND 2: CLAIMS 1, 5-9, 12-24 and 27-30 ARE UNPATENTABLE AS OBVIOUS OVER SHIEH IN VIEW OF SWENSON. 54

A. Claim 1 54
B. Claims 5-9 66
C. Claim 12 66
D. Claim 13 67
E. Claim 14 67
F. Claim 15 67
G. Claim 16 68
H. Claim 17 68
I. Claim 18 69
J. Claim 19 69
K. Claim 20 69
L. Claim 21 70
M. Claim 22 70
N. Claim 23 71
O. Claim 24 72
P. Claim 27 72
Q. Claims 28 72
R. Claims 29 73
S. Claim 30 73

XI. OBJECTIVE INDICIA OF NONOBVIOUSNESS 73

XII. DISCRETIONARY DENIAL UNDER § 325(D) OR § 314 IS NOT WARRANTED 74

XIII. MANDATORY NOTICES UNDER 37 C.F.R. § 42.8 77

A. Real Party-in-Interest 77

B. Related Matters 78

C. Lead and Back-up Counsel and Service Information 78

XIV. CONCLUSION.....78

APPENDIX A – CLAIM LISTING A-1

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 230 of 496

TABLE OF AUTHORITIES

Page(s)

Cases

Apple Inc. v. Fintiv, Inc.,
IPR2020-00019, Paper 11 (PTAB March 20, 2020) 75, 76

Dish Network LLC v. Broadband iTV, Inc.,
IPR2020-01359, Paper 15 (PTAB Feb. 12, 2021).....75

General Plastic Industrial Co., Ltd. v. Canon Kabushiki Kaisha,
PR2016-01357, Paper 19 (PTAB Sept. 6, 2017).....74

Orckit Corp. v. Cisco Systems, Inc.,
Case No. 2:22-cv-00276 (E.D. Tex.).....78

PEAG LLC v. Varta Microbattery GMBH,
IPR2020-01214, Paper 8 (Jan. 6, 2021).....76

Phillips v. AWH Corp.,
415 F.3d 1303 (Fed. Cir. 2015) (en banc) 5

Sand Revolution II LLC v. Continental Intermodal Group-Trucking LLC,
IPR2019-01393, Paper 24 (PTAB June 16, 2020) 75, 77

Sega of Am., Inc. v. Uniloc USA, Inc.,
IPR2014-01453, Paper 11 (PTAB Mar. 10, 2015).....73

Verizon v. Huawei,
IPR2020-01079, Paper 10 (Jan. 14, 2021).....77

Statutes

35 U.S.C. § 102(a) 4, 5

35 U.S.C. § 103.....4, 5

35 U.S.C. § 311.....1

35 U.S.C. §314(a) 1, 74

35 U.S.C. § 315(e)(2).....77

35 U.S.C. §325(d).....74

Other Authorities

37 C.F.R. § 42.8.....77

37 C.F.R. § 42.8(b)(1).....77

37 C.F.R. § 42.100.....1, 5

PETITIONERS' EXHIBIT LIST

1001	U.S. Patent No. 10,652,111 (the '111 Patent)
1002	Prosecution History of the '111 Patent
1003	Curriculum Vitae of Samrat Bhattacharjee
1004	Declaration of Dr. Samrat Bhattacharjee, dated February 21, 2023
1005	U.S. Patent No. 9,264,400 ("Lin")
1006	U.S. Patent Application Publication No. 2013/0291088 ("Shieh")
1007	U.S. Patent Application Publication No. 2013/0322242 ("Swenson)
1008	RFC 2460, Internet Protocol, Version 6 (IPv6) Specification
1009	Nunes, A., et al., A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks
1010	Complaint in *Orckit Corp. v. Cisco Systems, Inc.*, Case No. 2:22-cv-00276 (E.D. Tex.)
1011	Amended Scheduling Order in *Orckit Corp. v. Cisco Systems, Inc.*, Case No. 2:22-cv-00276 (E.D. Tex.)
1012	Interim Procedure for Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation, Vidal, K., United States Patent and Trademark Office, June 21, 2022
1013	Federal Case Management Statistics for the Eastern District of Texas, as of June 30, 2022

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 Claims 1-9, 12-24 and 27-31 (the "Challenged Claims") of U.S. 10,652,111 ("the '111 Patent," EX1001). The claimed methods in the '111 Patent would have been obvious to a person of ordinary skill in the art ("POSA") well before the '111 Patent's earliest priority date. For example, the combination of U.S. Patent No. 9,264,400 ("Lin," EX1005) and U.S. Patent Application Publication No. 2013/0333342 ("Swenson," EX1007) discloses deep packet inspection ("DPI") of packets in a computer network where a network node is under the control of a central controller. The combination of U.S. Patent Application Publication No. 2013/0291088 ("Shieh," EX1006) and Swenson teaches the same thing. The disclosures in these three prior art references, along with the knowledge of a POSA, render the Challenged Claims unpatentable as obvious, as explained below in Grounds 1 and 2.

II. GROUNDS FOR STANDING

Petitioner certifies that the '111 Patent is available for *inter partes* review and that Petitioner is not barred or estopped from requesting *inter partes* review challenging the patent claims on the grounds identified in this Petition.

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 234 of 496

III. REQUESTED RELIEF AND REASONS FOR REQUESTED RELIEF

Petitioner asks that the Board institute a trial for an *inter partes* review of the Challenged Claims, and that the Director cancel them as unpatentable. The analysis demonstrating the obviousness of the Challenged Claims is set forth in the below sections of this Petition and supported by the declaration of Petitioner's expert, Dr. Samrat Bhattacharjee. EX1004, ¶¶74-318; EX1003.

IV. OVERVIEW OF THE '111 PATENT

A. Summary of the '111 Patent

The '111 Patent discloses methods and systems relating to "deep packet inspection (DPI) in a software defined network (SDN)." EX1001, Abstract; *see id.*, 1:14-16, EX1004, ¶¶30-34. The '111 Patent discloses a "central controller of the SDN" that is used to "configure[e] a plurality of network nodes operable in the SDN" with instructions that tell the network nodes what to do with incoming packets. EX1001, 2:27-30, 2:3-3; EX1004, ¶30. For example, the central controller may send a "probe" instruction to a network node such that, when the network node receives a packet that matches a "packet-applicable criterion," the network node will "mirror" (i.e., send) some or all of the packet to a security component for inspection. EX1001, 2:3-44; EX1004, ¶30.

"[T]he central controller 111 [shown below in Figure 1 of the '111 Patent] is configured to perform deep packet inspection on designated packets from designated

flows or TCP sessions." EX1001, 4:5-7; EX1004, ¶31. "To this end, the central controller 111 is further configured to instruct each of the network nodes 112 which of the packets and/or sessions should be directed to the controller 111 for packet inspections." EX1001, 4:8-11. "The determination [of whether a packet requires inspection] is performed based on a set of instructions provided by the controller 111." EX1001, 4:14-15. "A packet that requires inspection is either redirected to the controller 111 or mirrored and a copy thereof is sent to the controller 111." EX1001, 4:15-18.

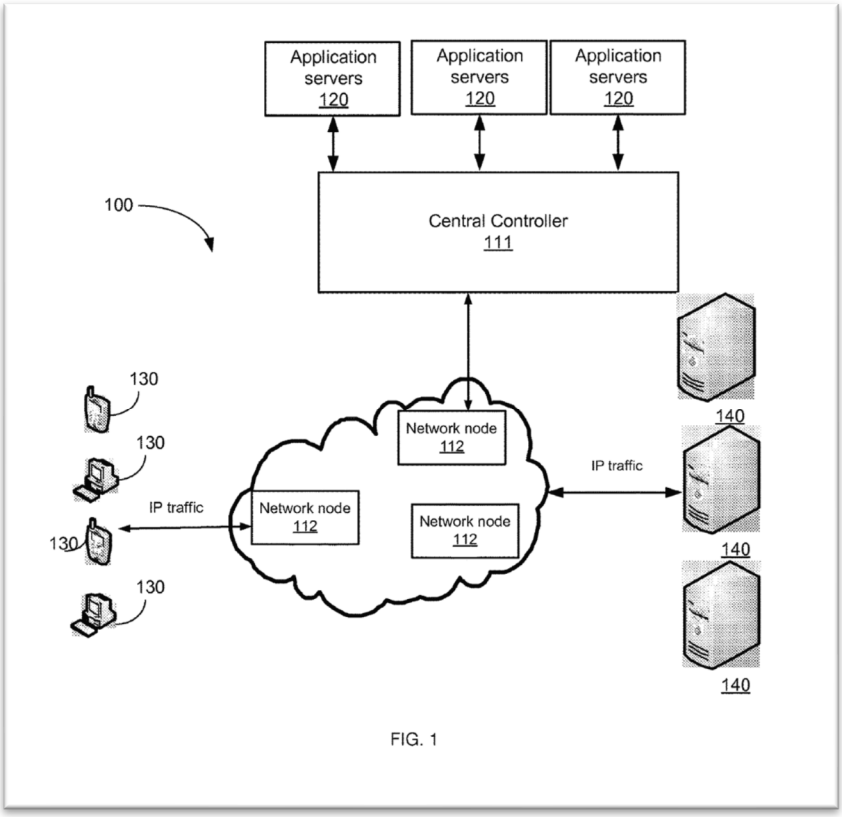


Figure 1 of the '111 Patent

During prosecution, the Applicant relied heavily on claim limitations reciting "… sending, by the controller to the network node over the packet network, an instruction and a packet-applicable criterion" and "… receiving by the network node from the controller, the instruction and the criterion" to distinguish the prior art, along with arguments that there was no motivation to combine the cited art. EX1002 at 322-330, 397-417, 492-501; EX1004, ¶¶35-46.

B. Priority Date

Solely for the purposes of this Petition, Petitioner assumes that the priority date for the '111 Patent is April 22, 2014, the filing date of U.S. Provisional Patent Application No. 61/982,358 to which the '111 Patent claims priority. EX1004, ¶47.

V. STATUTORY GROUNDS FOR CHALLENGES

Ground #1: Claims 1-9, 12-24 and 27-31 of the '111 Patent are obvious under 35 U.S.C. § 103 over Lin in view of Swenson and the knowledge of a POSA. Lin was filed on December 2, 2013, and issued on February 16, 2016. EX1005. Thus, Lin qualifies as prior art under at least post-AIA 35 U.S.C. § 102(a)(2).

Swenson claims priority to a pair of provisional applications filed on June 1, 2012, and January 18, 2013, respectively. Swenson was filed as a non-provisional application on May 31, 2013. Swenson published on December 5, 2013. Thus, Swenson qualifies as prior art under at least post-AIA 35 U.S.C. §§ 102(a)(1)-(2).

Ground #2: Claims 1, 5-9, 12-24 and 27-30 of the '111 Patent are obvious under 35 U.S.C. § 103(a) over Shieh in view of Swenson and the knowledge of a POSA. Shieh was filed as a provisional application on April 11, 2012 and as a non-provisional application on April 10, 2013. Shieh published on October 31, 2013. Thus, Shieh qualifies as prior art under at least post-AIA 35 U.S.C. §§ 102(a)(1)-(2). Swenson qualifies as prior art for the reasons stated above for Ground 1.

VI. LEVEL OF ORDINARY SKILL IN THE ART

As of April 22, 2014, a POSA would have had a bachelor's degree in computer science, computer engineering, or an equivalent, and two years of professional experience, and a POSA would have had a working knowledge of hardware and software for packet-switched networking. EX1004, ¶¶48-49. Lack of work experience can be remedied by additional education and vice versa. *Id.*, ¶48.

VII. CLAIM CONSTRUCTION

In *inter partes* review, claim terms must be given their ordinary and customary meaning as understood by a POSA at the time of the invention in light of the specification and the prosecution history pertaining to the patent. *See* 37 C.F.R. § 42.100(b); *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed. Cir. 2015) (en banc).

The claim term "controller" should be construed to mean "an entity configured to perform deep packet inspection on packets." EX1001, 10:52-62;

EX1004, ¶¶69-71. The '111 Patent discloses "a method for deep packet inspection (DPI) in a software defined network (SDN), wherein **the method is performed by a central controller** of the SDN." EX1001, 2:27-30 (emphasis added); *see id.*, 3:56-59. Further, the patent states that "the central controller 111 is **configured to perform deep packet inspection on designated packets** from designated flows or TCP sessions." *Id.*, 4:5-7 (emphasis added); *see id.*, 2:49-51, 4:8-11, 9:67-10:1.

Further, the '111 Patent describes that "the central controller 111 includes a DPI flow detection module 311, a DPI engine 312, and a memory 313, and a processing unit 314," as shown below in Figure 3. EX1001, 5:33-36. "The DPI engine 312 [is] configured to inspect a packet or a number of bytes to provide application metadata as required by an application executed by an application server 120." EX1001, 5:36-39; *see id.*, 5:40-59. A POSA would have known from this description that the central controller was configured in this manner to provide DPI on redirected packets, as all of the embodiments in the '111 Patent disclose that redirected packets are sent to the central controller for DPI. *See, e.g.*, EX1004, ¶71; EX1001, 4:8-18, 4:49-50, 8:1-5.

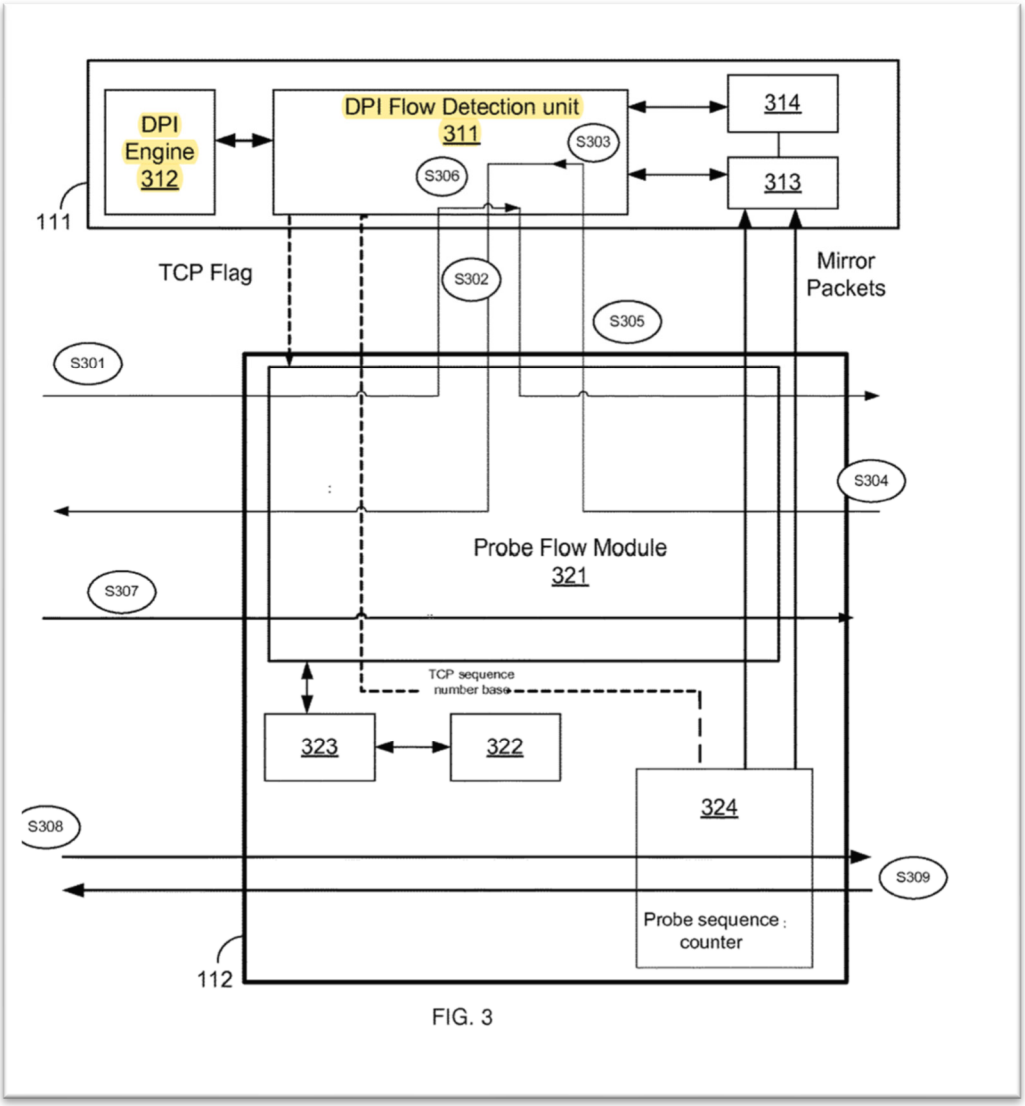


Figure 3 of the '111 Patent (Annotated)

EX1001, Figure 3; *see id.*, Figures 4-6.

Further, the claim term "instruction" should be construed to mean "a command to determine if a packet requires inspection or not." EX1001, 10:56-62; EX1004, ¶72. The '111 Patent discloses that "each network node 112 is configured to determine if an incoming packet requires inspection or not." EX1001, 4:12-14.

The patent states that "the central controller 111 is further configured to instruct each of the network nodes 112 which of the packets and/or sessions should be directed to the controller 111 for packet inspections." Moreover, the exemplary instructions provided in the '111 Patent are various commands used to determine whether or not a packet requires inspection. EX1001, 4:23-56; *see id.*, 8:23-32, 8:40-53, 9:26-28; EX1004, ¶72.

Terms not specifically construed have their plain and ordinary meaning as understood by a POSA. EX1004, ¶73.

VIII. THE ART PRIOR TO THE '111 PATENT

A. Lin

Lin "relates generally to computer security, and more particularly but not exclusively to software defined networking." EX1005, 1:7-9; *see id.*, Abstract; EX1004, ¶¶50-56.[1] "In one embodiment, a software defined networking (SDN) computer network includes an SDN controller and an SDN switch." EX1005, 1:58-60; Figures 6-8. "The SDN controller inserts flow rules in a flow table of the SDN switch to create an SDN pipe between a sender component and a security

[1] Background discussion of software defined networking can be found in Paragraphs 21-29 of Dr. Bhattacharjee's declaration and in EX1009.

component." EX1005, 1:60-62; *see id.*, 1:62-64, 4:8-31, 4:53-65, 6:1-12. "The SDN pipe allows outgoing packets sent by the sender component to be received by the security component." EX1005, 1:64-65; *see id.*, 3:25-31, 6:40-48. "The security component inspects the outgoing packets for compliance with security policies and allows the outgoing packets to be forwarded to their destination when the outgoing packets pass inspection." EX1005, 1:66-2:2; *see id.*, 3:31-33, 6:48-63, 7:9-21.

Figure 6 of Lin, reproduced below, shows "a schematic diagram of an SDN computer network 600" in which "[t]he SDN controller 610 provides a logically centralized framework for controlling the behavior of the SDN computer network 600." EX1005, 3:40-42, 4:7-9; *see id.*, 3:42-44, 4:9-12; EX1004, ¶51. "The SDN controller 610 may include a flow policy database 611." EX1005, 4:12-13. "The flow policy database 611 may comprise flow policies that are enforced by the controller 610 on network traffic transmitted over the SDN computer network 600." EX1005, 4:13-16; *see id.*, 4:16-18. "The flow policies may be enforced in terms of flow rules (labeled as 624) that are stored in the flow tables 621 of the SDN switch 620." EX1005, 4:18-20.

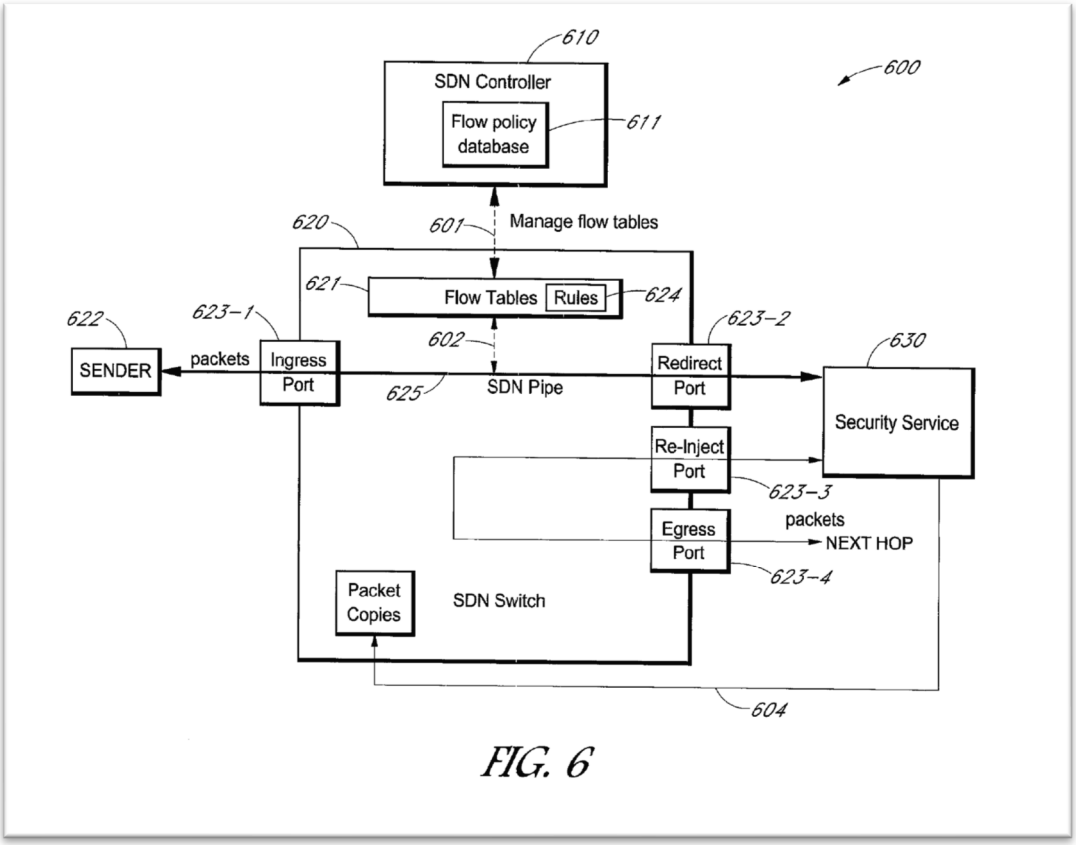


FIG. 6

Figure 6 of Lin

B. Shieh

Shieh relates "generally to network security" and discloses a "network system [that] includes a security device and a network access device." EX1006, ¶[0002]; EX1004, ¶¶57-62. "The network access device is to receive a packet from a source node destined to a destination node, and to examine a data structure maintained by the network access device to determine whether the data structure stores a data member having a predetermined value, the data member indicating whether the packet should undergo security processing," as shown below in Figure 1 of Shieh. EX1006, Abstract; *see id.*, ¶[0002], Figures 1, 2A, 3. "If the data member matches

the predetermined value, the packet is transmitted to a security device associated with the network access device to allow the security device to perform content inspection." *Id.*, Abstract; *see id.*, ¶[0042], ¶[0049]. "[I]n response to a response received from the security device, the packet is routed to the destination node dependent upon the response." *Id.*, Abstract; *see id.*, ¶[0017], ¶[0018], ¶[0023], ¶[0029], ¶[0037], Claim 1, Figure 2B; EX1004, ¶57.

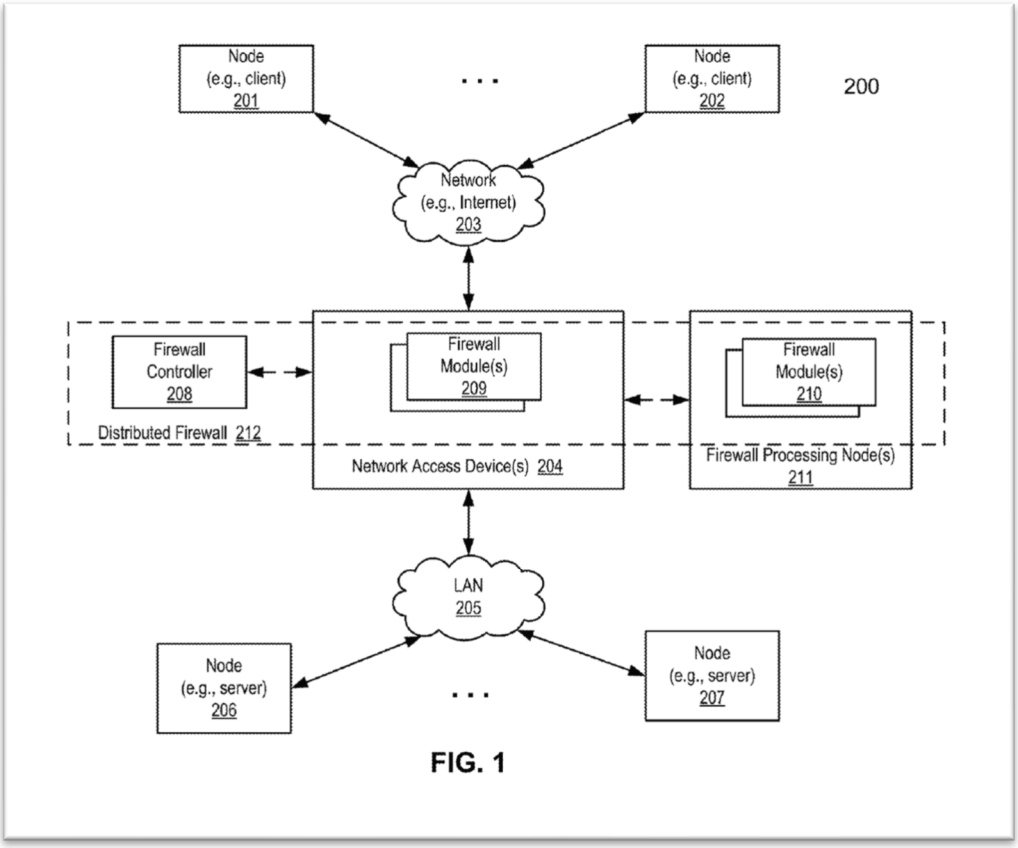


Figure 1 of Shieh

C. Swenson

Swenson discloses a system and method for "selectively monitoring traffic in a service provider network." EX1007, Abstract; *see id.*, ¶[0018]-¶[0022]; EX1004, ¶¶63-68. Figure 1 of Swenson (reproduced below) shows that "[t]he network 120 is a communication network that transmits data between the user devices 110, the steering devices 130 and the origin server 160 and/or the video optimizer 150." EX1007, ¶[0023]. "In one embodiment, the steering device 130 characterizes traffic routed through it to identify flows of interest for further inspection at the network controller 140." *Id.*, ¶[0026]; *see id.*, ¶[0058]. "Alternatively, the network controller 140 interfaces with the steering device 130 to coordinate the monitoring and characterization of network traffic, such as identifying large and small objects in HTTP traffic flows." *Id.*, ¶[0026]. "In this case, the steering device 130 receives instructions from the network controller 140 based on the desired criteria for characterizing flows of interest for further inspection." *Id.*, ¶[0026]. When a flow matches a particular signature, "the steering device 130 forwards the HTTP request and a portion of the HTTP response to the network controller 140 over the [Internet content adaption protocol] client interface 404." *Id.*, ¶[0059]; *see id.*, ¶[0060]. "After receiving the request and the portion of response at the ICAP server interface 406, the flow analyzer 312 of the network controller 140 performs a deep flow inspection

to determine if the flow is worth bandwidth monitoring and/or user detection." *Id.*, ¶[0059]; *see id.*, ¶[0060]; EX1004, ¶¶67-68.

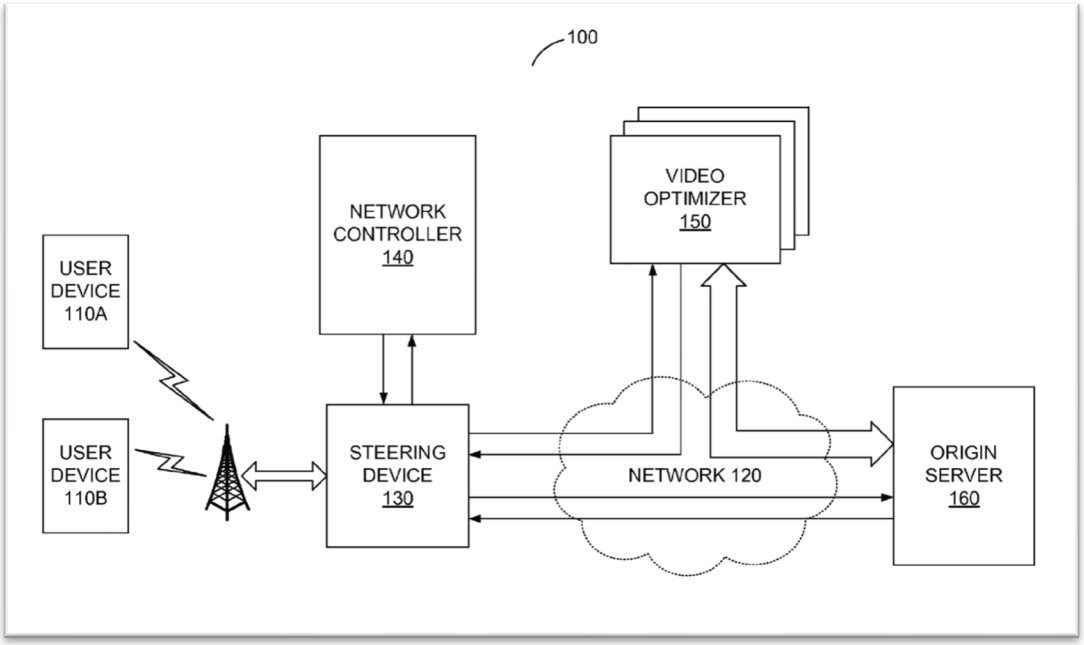


Figure 1 of Swenson

IX. GROUND 1: CLAIMS 1-9, 12-24 AND 27-31 ARE UNPATENTABLE AS OBVIOUS OVER LIN IN VIEW OF SWENSON.

The combination of Lin and Swenson, along with the knowledge of a POSA, renders Claims 1-9, 12-24 and 27-31 obvious. EX1004, ¶¶74-206.

A. Claim 1

[1.0] A method for use with a packet network including a network node for transporting packets between first and second entities under control of a controller that is external to the network node, the method comprising:

Element [1.0], to the extent it is limiting, is disclosed by Lin. EX1001, 10:51-55.

Lin discloses a method for use with a packet network. For example, Lin's Abstract states that it relates to "[a] software defined networking (SDN) **computer network**." EX1005, Abstract (emphasis added). The specification further states that "[t]he present invention relates generally to computer security, and more particularly but not exclusively to software defined networking." *Id.*, 1:7-9; *see id.*, 1:58-60, 2:47-65, 3:25-33, 3:40-64. A POSA would have known that this computer network is a packet network. EX1004, ¶¶75-76. Indeed, Lin refers to the "transmission of packets over the SDN computer network 600." EX1005, 4:19-21. Moreover, Figure 6 of Lin shows that the system disclosed in Lin is for a packet network, as can be seen below:

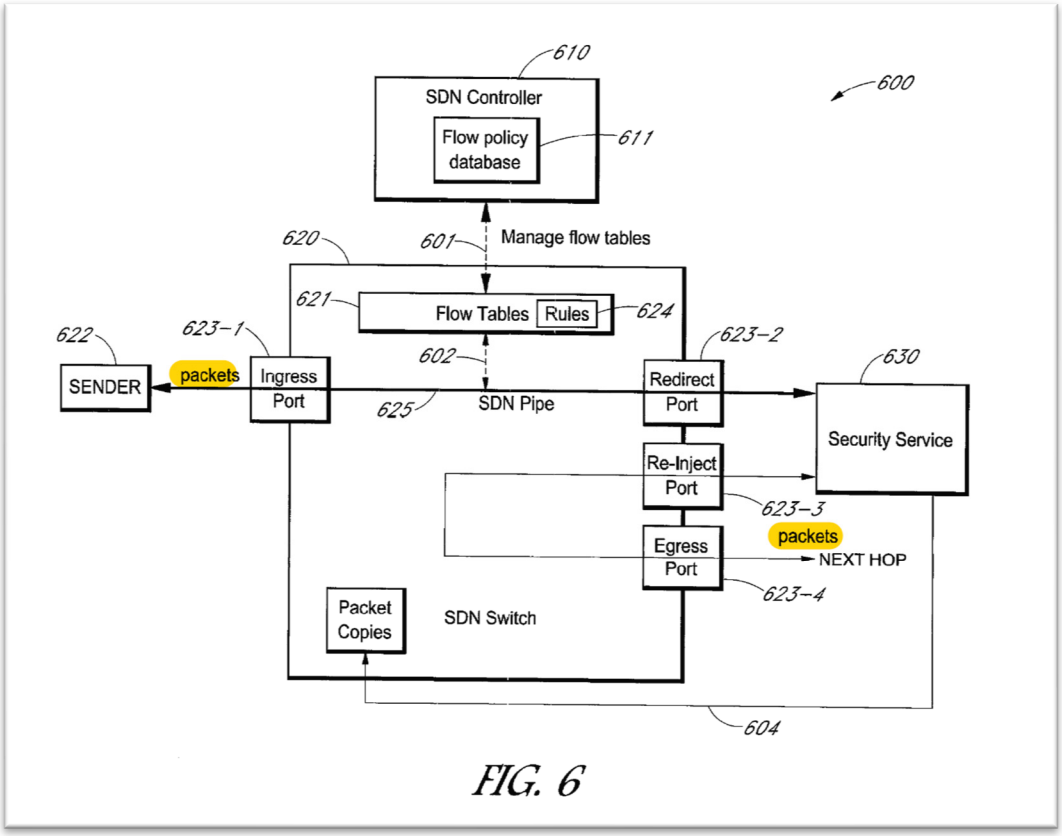


Figure 6 of Lin (Annotated)

EX1005, Figure 6; *see id.*, Figures 1-5, 7-9.

Further, Lin discloses an "SDN switch" that **corresponds to the claimed network node** for transporting packets between first and second entities. EX1005, 1:58-2:4, 4:33-67, 6:13-23, 6:57-63; EX1004, ¶77. Lin explains that the SDN switch transports packets from a "sender" component (**which corresponds to the claimed first entity**), through an ingress port, out an egress port, and to the "next hop" or destination (**which corresponds to the claimed second entity**), as shown below in Figure 6. EX1005, 1:58-2:4, 4:33-67, 6:13-23, 6:57-63, 7:10-23; 7:39-8:18, 9:63-10:22; EX1004, ¶77.

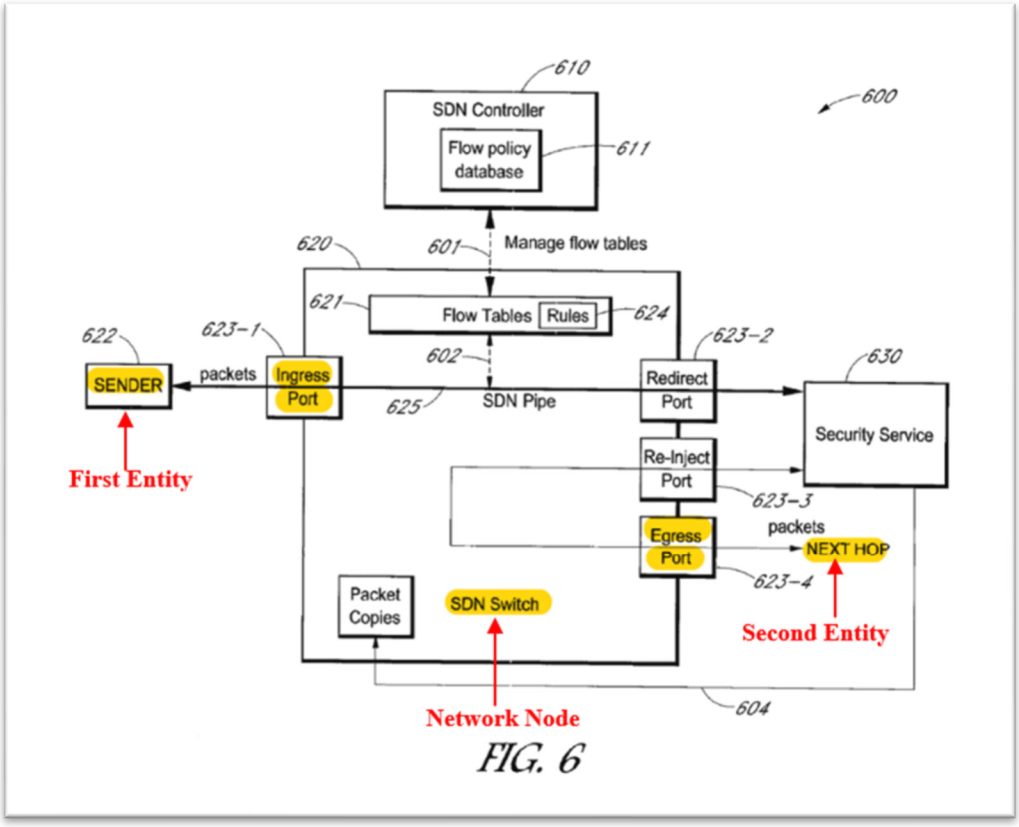


Figure 6 of Lin (Annotated)

EX1005, Figure 6.

Lin explains that the SDN switch is under the control of a "SDN controller" (**which corresponds to the claimed controller**) that is external to the SDN switch. EX1004, ¶78. Lin states that "the SDN controller 610 provides a logically centralized framework for controlling the behavior of the SDN computer network 600," including one or more SDN switches. EX1005, 4:8-31. The SDN controller includes a "flow policy database" that contains flow policies to control the transmission of packets through the SDN switch. *Id.*, 4:8-31; *see id.*, 1:58-2:4, 6:1-12. Lin explains that the SDN controller is external to the SDN switch (i.e., the

network node): "The SDN controller 610 and the SDN switch 620 are logically separate components." *Id.*, 3:51-52, *see id.*, 4:8-10. Further, Figures 1 and 6-8 of Lin show the SDN controller as external to the SDN switch:

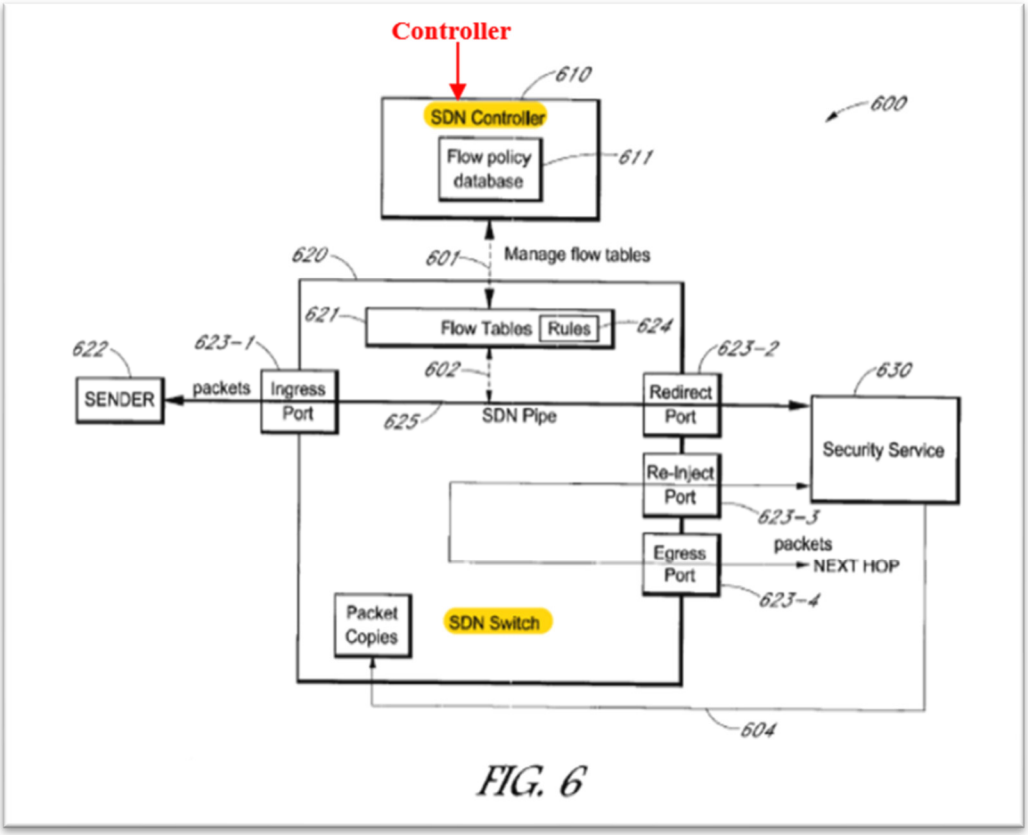


Figure 6 of Lin (Annotated)

Id., Figure 6; *see id.*, Figures 1, 7-8.

In addition, as addressed above in the claim construction section, the controller is an entity configured to perform DPI on packets.[2] EX1004, ¶79. The combination of Lin and Swenson discloses such a controller. Lin states that the analysis performed by security service 630 includes DPI: "Network security vendors provide network security services, such as firewall or deep packet inspection (DPI)." EX1005, 3:11-12. Moreover, Lin discloses that security service 630 "may also comprise a physical machine, e.g., a server computer, an appliance, or a gateway computer, etc." EX1005, 5:51-55. Further, Lin states that "[t]he security service 630 may be connected to the SDN switch 620 by a physical link (i.e., using a wire), a virtual link (i.e., in a virtualized environment), or by a software tunnel." A POSA would have known from these disclosures in Lin that the security service 630 can use the same hardware or software as the controller, and that the security service 630 can be connected to the SDN switch 620 in the same way as the controller. EX1004, ¶79. Thus, a POSA would have understood that one of the limited number of design options would have been to implement the security service as part of a controller configured to perform DPI analysis on packets, and a POSA would have had a

[2] To the extent that the PTAB does not agree with this construction, Lin still discloses Element [1.0] for the reasons discussed above.

reasonable expectation that the controller would have been successful in performing DPI analysis. EX1004, ¶79.

Further, Swenson teaches the use of a controller configured to perform DPI. EX1004, ¶80. Swenson discloses that, when its system detects a HTTP packet flow matching a particular signature, "the steering device 130 forwards the HTTP request and a portion of the HTTP response to the network controller 140 over the ICAP client interface 404." EX1007, ¶[0059]; *see id.*, ¶[0060]. "After receiving the request and the portion of response at the ICAP server interface 406, **the flow analyzer 312 of the network controller 140 performs a deep flow inspection** to determine if the flow is worth bandwidth monitoring and/or user detection." EX1007, ¶[0059] (emphasis added); *see id.*, ¶[0060] (stating that the "controller 140 ingests the network flow for inspection"), Figures 1, 4A-4B. A POSA would have known that a "flow" is a series of packets having a specific signature. EX1004, ¶80. As such, it would have been obvious to a POSA that Swenson's reference to "deep flow inspection" refers to performing DPI on one or more packets in a flow. *Id.*

In addition, Swenson discusses "an example event trace of [Swenson's] 'continue' working mode" in which the steering device 130 "sends an ICAP request message 516 comprising [a] HTTP GET request header and a portion of the [HTTP] response payload to the network controller 140, which inspects the message to

determine whether to monitor the flow or optimize the video." EX1007, ¶[0065], Figure 5. A POSA would have known that the analysis of the "response payload" to be DPI by a DPI-capable controller, as DPI refers to monitoring the payload of a packet. EX1004, ¶81. Swenson thus demonstrates that it would have been well-known to a POSA as of the priority date for the '111 Patent to implement a controller configured to perform DPI analysis in a system such as Lin. EX1007, ¶[0046], ¶[0059]-¶[0060], ¶[0073], ¶[0076]-[0077], ¶[0084]-¶[0086], Figures 1 and 4A-4B; EX1004, ¶81.

A POSA would have been motivated to combine the teachings of Lin and Swenson. EX1004, ¶¶82-87. Lin and Swenson are analogous art references that address the same technology and attempt to resolve the same issues relating to routing network traffic in an efficient manner that conserves network resources. EX1005, 1:58-2:4, 4:8-31, 5:8-55, 6:1-12, 6:40-63, 7:24-8:18, Figures 6-9; EX1007, ¶[0023]-¶[0032], ¶[0038]-¶[0043], ¶[0057]-¶[0061], Figures 1-4A; EX1004, ¶82. They each use a central controller to provide instructions and packet-applicable criterion to network nodes to determine which packets should be redirected and which packets can be sent directly to a destination node. EX1005, 1:58-2:4, 4:8-31, 5:8-55, 6:1-12, 6:40-63, 7:24-8:18, Figures 6-9; EX1007, ¶[0023]-¶[0032], ¶[0038]-¶[0043], ¶[0057]-¶[0061], Figures 1-4A; EX1004, ¶82. Lin provides a method in which a "SDN controller inserts flow rules in a flow table of the SDN switch" that

provide instructions and packet-applicable criterion to the SDN switch for use in identifying which packets should be sent to a security service. EX1005, 1:60-2:4, 3:21-24, 4:8-31, 4:53-67, 6:1-12, 6:54-63, Figure 6. Lin also discloses that the security service sends packets to their destination if the packets pass inspection. *Id.* Further, Lin explains that "bypass rules [that] are inserted in the flow tables 621 such that particular packets that do not need to be inspected are not redirected to the security service 630." EX1005, 7:23-27; *see id.*, 9:12-16, Figure 9.

Similarly, Swenson discloses a method in which a network controller sends instructions to steering devices through which packet flows pass. EX1007, EX1007, ¶[0023]-¶[0032], ¶[0038]-¶[0043], ¶[0057]-¶[0061], Figures 1-4A; EX1004, ¶83. Swenson states that, when its systems detect a HTTP packet flow matching a particular signature, "the steering device 130 forwards the HTTP request and a portion of the HTTP response to the network controller 140 over the ICAP client interface 404." EX1007, ¶[0059]; *see id.*, ¶[0060]. Like Lin, Swenson teaches that only certain packets are redirected to the controller in order to be efficient in analyzing packets of interest. EX1007, ¶[0059]-¶[0060]; EX1004, ¶83. Further, a POSA would have understood that Swenson, like Lin, could involve analyzing packets for a security function. EX1004, ¶83. A POSA would have known that the bandwidth monitoring in Swenson can be used as a security application that monitors for Denial of Service ("DOS") attacks that occupy significant bandwidth in a

network. EX1007, ¶[0059]-¶[0060]; *see id.*, ¶[0039] (explaining that the controller can incorporate "security functions"); EX1004, ¶83. Similar to Lin, Swenson teaches that packets can be routed or communicated to a destination node after the processing of those packets is completed. EX1007, ¶[0060]. Moreover, the architecture of Swenson is substantially similar to the architecture of Lin, as shown in the below figures showing a central controller directly connected to a network node (the SDN switch in Lin and the steering device in Swenson) through which packet flows pass:

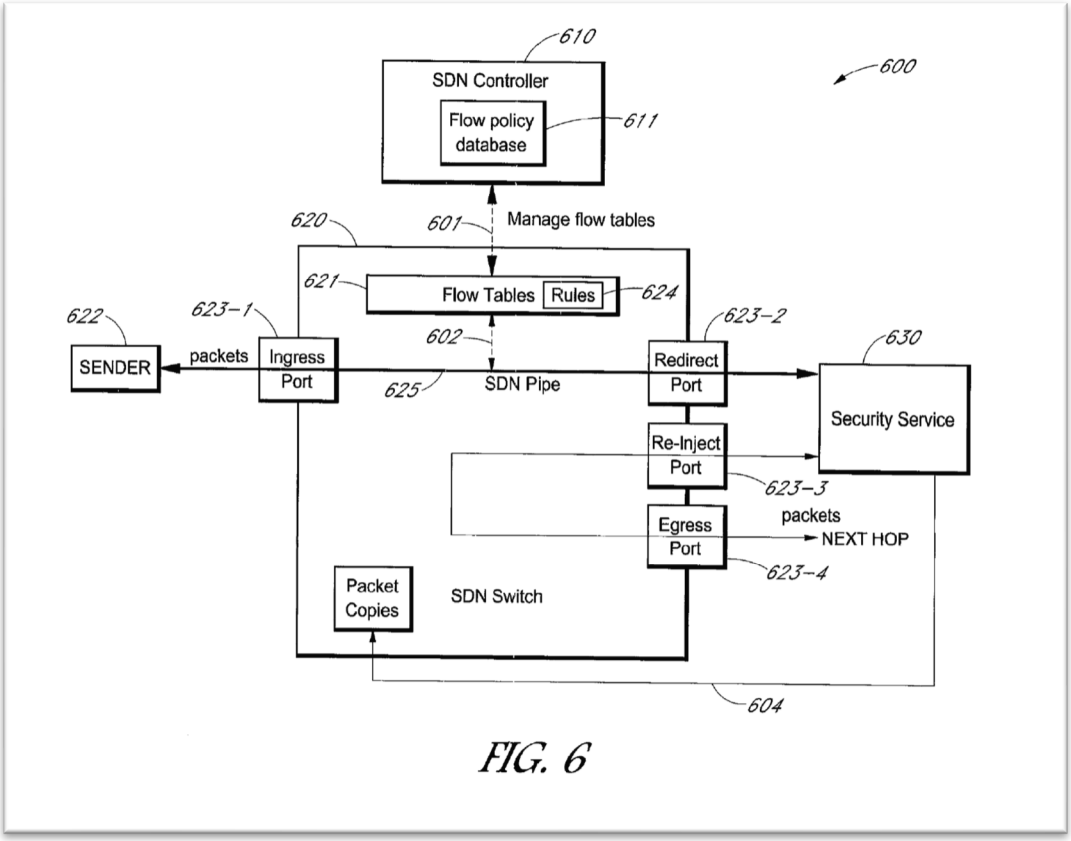


FIG. 6

Figure 6 of Lin

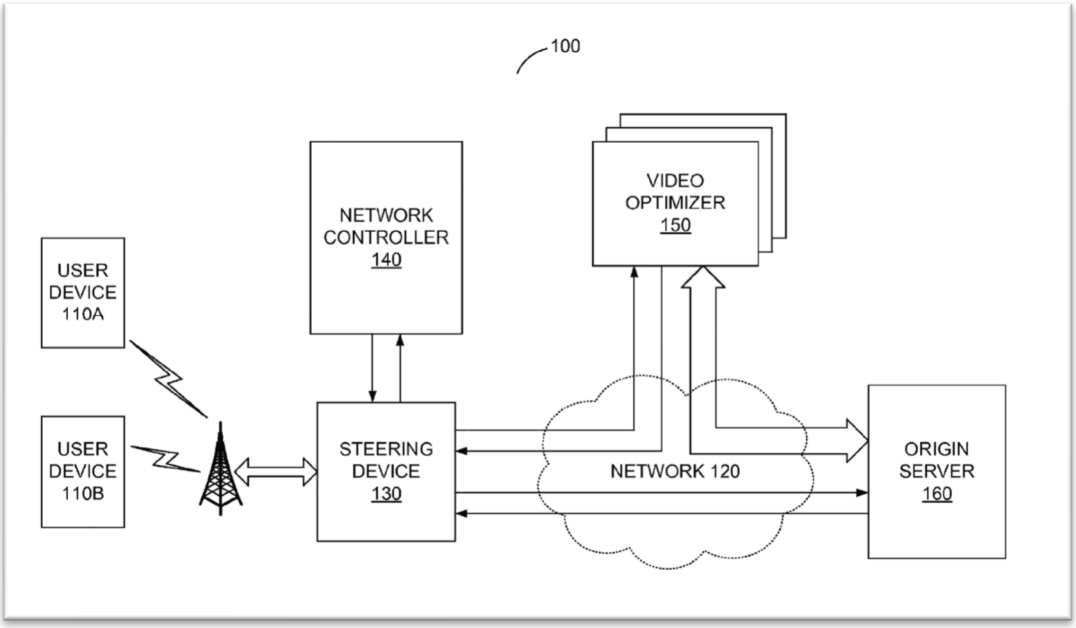


Figure 1 of Swenson

EX1004, ¶83.

The determination of which packets to redirect (as opposed to redirecting all packets) increases the efficient operation of a computer network. EX1004, ¶84. Lin and Swenson each teach that only certain packets may need to be redirected for further processing before being sent to the destination node. EX1005, 1:58-2:4, 4:8-31, 5:8-55, 6:1-12, 6:40-63, 7:24-8:18, Figures 6-9; EX1007, ¶[0023]-¶[0032], ¶[0038]-¶[0043], ¶[0057]-¶[0061], Figures 1-4A. This allows the methods in Lin and Swenson to improve the operation of transporting packets across a network while maintaining a proper speed for processing those packets by allowing the system to stop redirecting certain packets in a flow if it is no longer necessary to do so. EX1004, ¶84.

Further, a POSA would have found it obvious to implement the security processing module in Lin as part of the controller in light of the disclosures in Swenson to send packets to the controller for DPI. EX1007, ¶[0059]-¶[0060], Figures 1, 4A-4B; EX1004, ¶85. A POSA would have understood there were efficiencies to implementing Lin's security processing module as part of the controller. EX1004, ¶85. For example, the use of a central location for packets from different nodes to undergo inspection by a security component allows the same security algorithms to be applied to each analyzed packet. Further, a POSA would have known that security algorithms in a security component are often updated via software updates as new threats are identified. *Id.* A POSA would have understood that an efficient way to keep the security component up-to-date would be to have a central security component that is part of the central processor. *Id.*

A POSA would have had a reasonable expectation of success in modifying Lin to implement the disclosure of Swenson to route packets to the central controller for inspection. EX1004, ¶86. As discussed above, Lin and Swenson have a similar architecture in which network nodes are controlled by a separate external controller. EX1005, 1:58-2:4, 4:8-31, 5:8-55, 6:1-12, 6:40-63, 7:24-8:18, Figures 6-9; EX1007, ¶[0023]-¶[0032], ¶[0038]-¶[0043], ¶[0057]-¶[0061], Figures 1-4A. If Swenson could be constructed to have the network nodes send packets to the central controller for security inspection, a POSA would have understood that the same arrangement

could be implemented in Lin. EX1004, ¶86. Indeed, it would have been common sense for a POSA to modify Lin in this way after reading Swenson. *Id.*

[1.1] sending, by the controller to the network node over the packet network, an instruction and a packet-applicable criterion;

Element [1.1] is disclosed by Lin. EX1001, 10:56-58.

Lin explains that a controller (i.e., the SDN controller) controls network nodes (i.e., the SDN switches) in a packet network for the reasons stated above for Element [1.0]. EX1005, 1:58-2:4, 3:25-64; 4:8-31, 4:53-67, 6:1-12, 6:40-64, 7:39-8:18; Figures 6-9; EX1004, ¶¶88-89.

Lin also describes an instruction sent by the controller to the network node over the packet network that includes a command to determine whether or not a packet requires inspection. EX1004, ¶90. Specifically, the SDN controller inserts "flow rules in a flow table of the SDN switch to create an SDN pipe between a sender component and a security component." EX1005, 1:58-2:4. **These flow rules correspond to the claimed instruction**, and they are sent to the SDN switches to provide instruction to the SDN switches on what packets should be sent to the security component for inspection. *Id.*; *see id.*, 4:14-31, 4:53-67, 6:1-12; 7:39-8:18.

Further, Lin explains that a separate packet-applicable criterion sent by the controller to the network node over the packet network. EX1004, ¶91. Lin describes that the flow rules "may indicate inspection of particular packets (e.g., those that

meet one or more conditions) by a security service 630." EX1005, 4:21-31; *see id.*, 1:60-62, 5:8-12. A POSA would have known that the "one or more conditions" used to identify "particular packets" for inspection **would be packet-applicable criterion**. EX1004, ¶¶91-92; EX1005, 4:23-31, 5:8-12.

Indeed, Lin discloses examples of the criterion applicable to a particular packet, including "'IN_PORT', 'MAC src' (media access control (MAC) address of the source of the packet), 'MAC dst' (MAC address of the destination of the packet), 'IP src' (Internet Protocol (IP) address of the source of the packet), 'IP dst' (IP address of the destination of the packet)." EX1005, 5:16-21; EX1004, ¶91. Lin states, "When the conditions are met, i.e., the particular packet is identified, the action indicated in the corresponding 'Action' column is performed on the packet." EX1005, 5:22-24; *see id.*, 5:26-36, 6:1-12, 6:40-54, Table 1. Further, in discussing the examples shown in Table 2 and Table 3, Lin discloses "bypass flow rules" where packet-applicable criterion are used to identify HTTP packets. *Id.*, 2:2-4, 7:24-8:18.

Further, Lin explains that the flow rules containing the packet-applicable criterion are sent by the SDN controller to the flow table in the SDN switch. EX1004, ¶93. Lin states "[t]he SDN controller 610 may insert flow rules in the flow tables 621 (see arrow 601) to create an SDN pipe (labeled as 625) between the sender component 622 and the security service 630," as shown below in Figure 6. EX1005, 6:1-4; *see id.*, 6:40-41. With respect to Figure 6, Lin discloses an embodiment in

which "bypass flow rules are inserted in the flow tables 621 such that particular packets that do not need to be inspected are not redirected to the security service 630." *Id.*, 7:24-26; *see id.*, 7:27-8:18, Tables 2-3. A POSA would have understood from these disclosures that Lin discloses a packet-applicable criterion sent by the controller to the network node (i.e., the SDN switch) over the packet network. EX1004, ¶¶93-94.

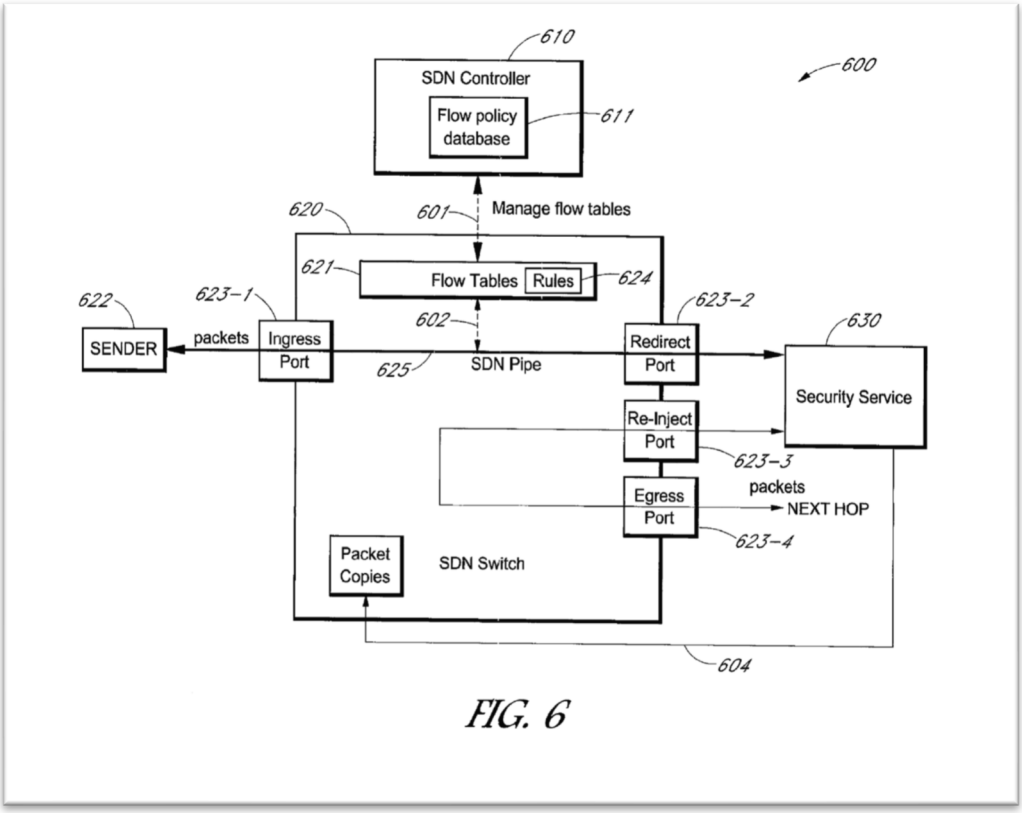


FIG. 6

Figure 6 of Lin

[1.2] receiving, by the network node from the controller, the instruction and the criterion;

Element [1.2] recites that the instruction and packet-applicable criterion that are sent by the controller to the network node in Element [1.1] are received by the

network node. EX1001, 10:59-60. As such, Lin discloses Element [1.2] for substantially the same reasons that Lin discloses Element [1.1]. EX1004, ¶¶95-100.

[1.3] receiving, by the network node from the first entity over the packet network, a packet addressed to the second entity;

Element [1.3] is disclosed by Lin. EX1001, 10:60-62.

Lin discloses that the SDN switch (i.e., the network node) receives packets from a "sender" component (which corresponds to the claimed first entity) and ultimately sends at least some packets to the "next hop" or destination (which corresponds to the claimed second entity), as shown below in Figure 6. EX1005, 1:58-2:4, 4:33-67, 6:13-23, 6:57-63, 7:10-23; 7:39-8:18, 9:63-10:22; EX1004, ¶¶101-103. A POSA would have understood that the packet is addressed to the next hop or second entity. EX1004, ¶102. The IP address of the destination of the packet is one of the criterion that is specifically discussed in Lin. EX1005, 5:16-31.

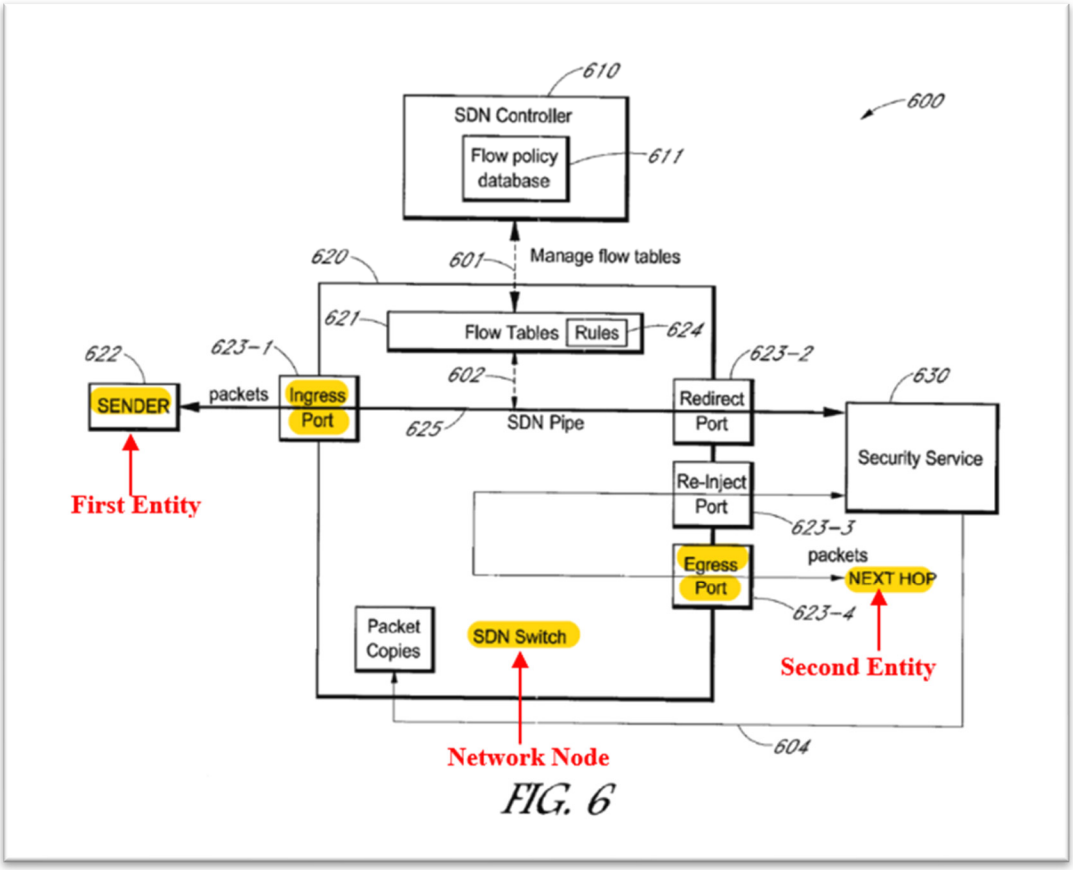


Figure 6 of Lin (Annotated)

Id., Figure 6.

[1.4] checking, by the network node, if the packet satisfies the criterion;

Element [1.4] is disclosed by Lin. EX1001, 10:63-64.

As discussed above for Element [1.1], Lin discloses that its SDN switch implements flow rules that "may indicate inspection of particular packets (e.g., those that meet one or more conditions) by a security service 630." EX1005, 4:23-31; *see id.*, 5:8-12, 6:1-4, 6:40-41, 7:24-8:18. A POSA would have understood this disclosure to mean that the "one or more conditions" used to identify "particular

packets" for inspection would be packet-applicable criterion, and that the SDN switch checks to determine whether incoming packets satisfy the packet-applicable criterion. EX1004, ¶104-105; EX1005, 4:23-31, 5:8-12. Indeed, Lin provides examples of the conditions (or criterion) applicable to a particular packet that are checked by the SDN switch, including "IN_PORT," "MAC src," "MAC dst," "IP src," and "IP dst." EX1005, 5:16-21; EX1004, ¶105. Lin states, "When the conditions are met, i.e., the particular packet is identified, the action indicated in the corresponding 'Action' column is performed on the packet." EX1005, 5:22-24; *see id.*, 5:26-36, 6:1-12, 6:40-54, Table 1. Further, in discussing the examples shown in Table 2 and Table 3, Lin discloses that the SDN switch implements "bypass flow rules" that check whether the packet meets certain packet-applicable criterion, such as identification of HTTP packets, which indicate that the packet should be routed to the destination node instead of the security device. *Id.*, 7:24-8:18; EX1004, ¶¶105-106.

[1.5] responsive to the packet not satisfying the criterion, sending, by the network node over the packet network, the packet to the second entity; and

Element [1.4] is disclosed by Lin. EX1001, 10:65-67.

Lin explains that the SDN switch implements "bypass flow rules" that check whether a packet meets certain packet-applicable criterion. EX1005, 7:24-8:18; *see id.*, 5:16-24, 5:26-36, 6:1-12, 6:40-54; EX1004, ¶¶107-109. In certain embodiments

of the bypass flow rules, if the packet does not satisfy the packet-applicable criterion, the packet is routed to the destination node (i.e., the second entity) instead of the security device. EX1005, 7:24-8:18. For example, Lin's discussion of Table 3 teaches that the SDN switch checks for a specific packet-applicable criterion – that port 80 is the source or destination port of the packets, which indicates that they are HTTP packets. *Id.*, 7:64-8:18. Lin states that, if a packet does not satisfy this criterion (i.e., the packet does not indicate that port 80 is the source or destination port), then the SDN switch (which corresponds to the claimed network node) sends the packet over the packet network to its destination node (which corresponds to the claimed second entity). *Id.*, 8:10-18; EX1004, ¶¶109-110. Specifically, with respect to Table 3 (reproduced below), Lin explains that "the bottom two rows [highlighted below] are bypass flow rules" that cause non-HTTP packets to "bypass the SDN pipe."

TABLE 3

IN_PORT	. . .	IP src	TCP src port	TCP dst port	. . .	Action	Count
Ingress_port_ID	*	*	*	80	*	Redirect port	10
Redirect_port_ID	*	*	80	*	*	Ingress port	10
Ingress_port_ID	*	*	*	*	*	Egress port	130
Egress_port_ID	*	*	*	*	*	Ingress port	130

Table 3 of Lin (Annotated)

EX1005, 8:10-18, Table 3.

[1.6] responsive to the packet satisfying the criterion, sending the packet, by the network node over the packet network, to an entity that is included in the instruction and is other than the second entity.

Element [1.6] is disclosed by Lin. EX1001, 11:1-4.

As discussed above for Element [1.5], Lin discloses that the SDN switch implements "bypass flow rules" that check whether a packet meets certain packet-applicable criterion. EX1005, 7:24-8:18; *see id.*, 5:16-24, 5:26-36, 6:1-12, 6:40-54; EX1004, ¶¶111-114. Further, in some embodiments, the bypass flow rules provide that packet which meet certain packet-applicable criterion should be routed to the security device. EX1005, 7:24-8:18. For example, Lin's discussion of Table 3 teaches that the SDN switch checks for a specific packet-applicable criterion – that port 80 is the source or destination port of the packets, which indicates that they are HTTP packets. *Id.*, 7:64-8:18. Lin states that, if a packet satisfies this criterion (i.e., the packet indicates that port 80 is the source or destination port), then the SDN switch (which corresponds to the claimed network node) sends the packet over the packet network to the security device (which corresponds to the claimed entity that is included in the instruction and is other than the second entity). *Id.*, 8:10-18, Table 3 (reproduced below.)

TABLE 3

IN_PORT	. . .	IP src	TCP src port	TCP dst port	. . .	Action	Count
Ingress_port_ID	*	*	*	80	*	Redirect port	10
Redirect_port_ID	*	*	80	*	*	Ingress port	10
Ingress_port_ID	*	*	*	*	*	Egress port	130
Egress_port_ID	*	*	*	*	*	Ingress port	130

Table 3 of Lin (Annotated)

Id., 8:10-18 (emphasis added), Table 3; EX1004, ¶113.

B. Claim 2

[2] The method according to claim 1, wherein the instruction is 'probe', 'mirror', or 'terminate' instruction, and upon receiving by the network node the 'terminate' instruction, the method further comprising blocking, by the network node, the packet from being sent to the second entity and to the controller.

Claim 2 is obvious in view of Lin and the knowledge of a POSA. EX1001, 11:5-10.

Lin teaches that the instruction can be a "terminate" instruction. EX1004, ¶¶115-116. A POSA would have understood that a terminate instruction was a well-known instruction that drops packets from the network such that they are no longer forwarded to a destination within the network. *Id.* Lin discloses that one of its "packet manipulation actions" is "dropping the packet." EX1005, 1:28-32. Lin also states that the security service 630 "may instruct the SDN switch 620 to drop the copied packets…" *Id.*, 7:19-22. A POSA would have understood these passages

from Lin to teach that the flow rules in Lin could include a terminate instruction. EX1004, ¶116.

Where the SDN switch in Lin receives a terminate instruction, a POSA would have understood that the terminate instruction easily could be implemented such that the SDN switch blocks a packet from being sent to the destination node (i.e., the claimed second entity) and the controller. EX1004, ¶117. Lin states that the SDN switch can be instructed to drop (i.e., block) packets before the packets are sent to their destination node. EX1005, 1:28-32, 7:19-22. It would have been straightforward for the flow rules in Lin to also instruct that packets are dropped by the SDN switch before they can be sent to controller. EX1004, ¶¶117-118. For example, a POSA would have known to implement a terminate instruction to drop (i.e., block) packets received at the SDN switch from an IP address that is known to originate malicious code, and thereby prevent those packets from being sent to the controller (or destination node) where the malicious code could do damage. *Id.*

C. Claim 3

[3] The method according to claim 1, wherein the instruction is a 'probe', a 'mirror', or a 'terminate' instruction, and upon receiving by the network node the 'mirror' instruction and responsive to the packet satisfying the criterion, the method further comprising sending the packet, by the network node, to the second entity and to the controller.

Claim 3 is obvious in view of Lin. EX1001, 11:11-16.

Further, Lin teaches that the instruction can be a "mirror" instruction. EX1004, ¶¶119-120. A POSA would have understood that a mirror instruction was a common instruction that can operate in one of two ways: (1) redirecting a packet from its intended destination to a new destination (such as a security service) or (2) making a copy of a packet and sending the copy of the packet to a new destination (such as a security service). *Id.* Lin discloses both methods of mirroring a packet. For example, Lin states that incoming packets can be "redirected or mirrored to the security service." EX1005, 3:25-33. Lin also discloses that a copy of a packet can be made, the packet mirrored to the security service, and an action (such as drop, forward or quarantine) taken on the copied patent. *Id.*, 7:10-22. A POSA would have understood these passages from Lin to teach that the flow rules in Lin could include a mirror instruction. EX1004, ¶120.

Where the SDN switch in Lin receives a mirror instruction, and the packet satisfies the criterion (as discussed above for Element [1.6]), Lin teaches that the mirror instruction instructs that "the SDN switch 620 may be configured to copy packets that are redirected to the security service 630 for inspection" and to "forward the copied packets to their destinations." EX1005, 7:10-22. As discussed above for Element [1.6], it would have been obvious to implement the security service 630 in Lin as part of the controller; thus, the original packet would be sent to the controller

by the SDN switch and a copy of the packet is sent to the destination node. *Id.*; EX1004, ¶¶121-123.

D. Claim 4

[4] The method according to claim 1, wherein the instruction is 'probe', 'mirror', or 'terminate' instruction, and upon receiving by the network node the 'probe' instruction and responsive to the packet satisfying the criterion, the method further comprising: sending the packet, by the network node, to the controller; responsive to receiving the packet, analyzing the packet, by the controller; sending the packet, by the controller, to the network node; and responsive to receiving the packet, sending the packet, by the network node, to the second entity.

Claim 4 is obvious over Lin, in view of the knowledge of a POSA. EX1001, 11:17-26.

Lin teaches that the instruction can be a "probe" instruction. EX1004, ¶¶124-125. A POSA would have understood that a probe instruction was a well-known instruction to send an incoming packet to a security service so that the packet can be probed or inspected for malicious code. *Id.* As discussed above for Element [1.1], Lin describes the SDN switch (i.e., the network node) receiving an instruction in the form of a flow rule that instructs the SDN switches that packets should be sent to the security component for analysis. EX1005, 1:58-2:4; *see id.*, 4:14-31, 4:53-67, 6:1-12; 9:23-40, Figures 6-9. Lin teaches an example wherein the flow rules received by the SDN switch instruct the SDN switch to redirect non-HTTP packets to the security service 630 for inspection. *Id.*, 7:39-63; *see id.*, 7:64-8:18 (instructing the

SDN switch to redirect "HTTP packets to the security service 630 for inspection"). A POSA would have understood from these disclosures in Lin that the flow rules sent to the SDN switch include a probe instruction for sending the packets to the security component. EX1004, ¶125.

Where the SDN switch in Lin receives a probe instruction, and the packet satisfies the criterion (as discussed above for Element [1.6]), Lin teaches that the probe instruction instructs the SDN switch to send the packet to the security service 630 for inspection. EX1005, Abstract, 1:58-2:4, 4:53-67, 5:26-36, 6:1-12, 6:40-63, 7:10-8:18, Figures 6-9. It would have been obvious in view of Swenson to implement the security service 630 in the controller (as discussed above for Element [1.0]); thus, the packet is sent to the controller for inspection. EX1004, ¶126.

Further, a POSA would have known from Lin that, when the packet is received by the controller, the packet is analyzed by the security service in the controller. EX1004, ¶127; EX1005, 6:48-50, 5:45-55, 7:10-22, 8:33-35, Figure 9. A POSA also would have known that, when the analysis is complete, the packet is sent from the controller back to the SDN switch and on to the destination node (i.e., the second entity). EX1004, ¶127; EX1005, 6:54-63, 7:10-22, 8:35-45, Figure 9.

E. Claim 5

[5] The method according to claim 1, further comprising responsive to the packet satisfying the criterion and to the instruction, sending the packet or a portion thereof, by the network node, to the controller.

Claim 5 is obvious in view of the combination of Lin and Swenson. EX1001, 11:27-30.

As discussed above for Claim 1, Lin discloses that a network node has an instruction and packet-applicable criterion and sending the packet to the controller responsive to the instruction and to the packet satisfying the packet-applicable criterion. EX1004, ¶¶129-130.

Moreover, Swenson teaches that the network node sends a packet or a portion thereof to the controller in response to the packet satisfying a criterion. EX1004, ¶131. As discussed above for Element [1.0], Swenson discloses that, when its system detects a HTTP packet flow matching a particular signature, "the steering device 130 **forwards the HTTP request and a portion of the HTTP response to the network controller 140** over the ICAP client interface 404." EX1007, ¶[0059] (emphasis added]; *see id.*, ¶[0060], ¶[0065], Figures 1, 4A-4B. A POSA would have understood this to mean that the steering device 130 corresponds to at least part of the claimed network node, and the signature that is compared to the HTTP packet flow corresponds to the claimed criterion. EX1004, ¶¶131-132. Responsive to the HTTP packet flow satisfying that criterion, the steering device applies an instruction pursuant to which the HTTP request and a portion of the HTTP response is sent to the network controller 140. EX1007, ¶[0059].

F. Claim 6

[6] The method according to claim 5, further comprising storing the received packet or a portion thereof, by the controller, in a memory.

Claim 6 is obvious over Swenson in view of the knowledge of a POSA. EX1001, 11:31-33.

Swenson discloses a flow cache memory 322 that is part of the network controller 140, as shown below in Figure 3 of Swenson. EX1004, ¶¶133-134. Further, Swenson states that "[o]nce a flow is reported to the network controller 140, a flow cache entry is created for the flow in the flow cache 322." EX1007, ¶[0061].

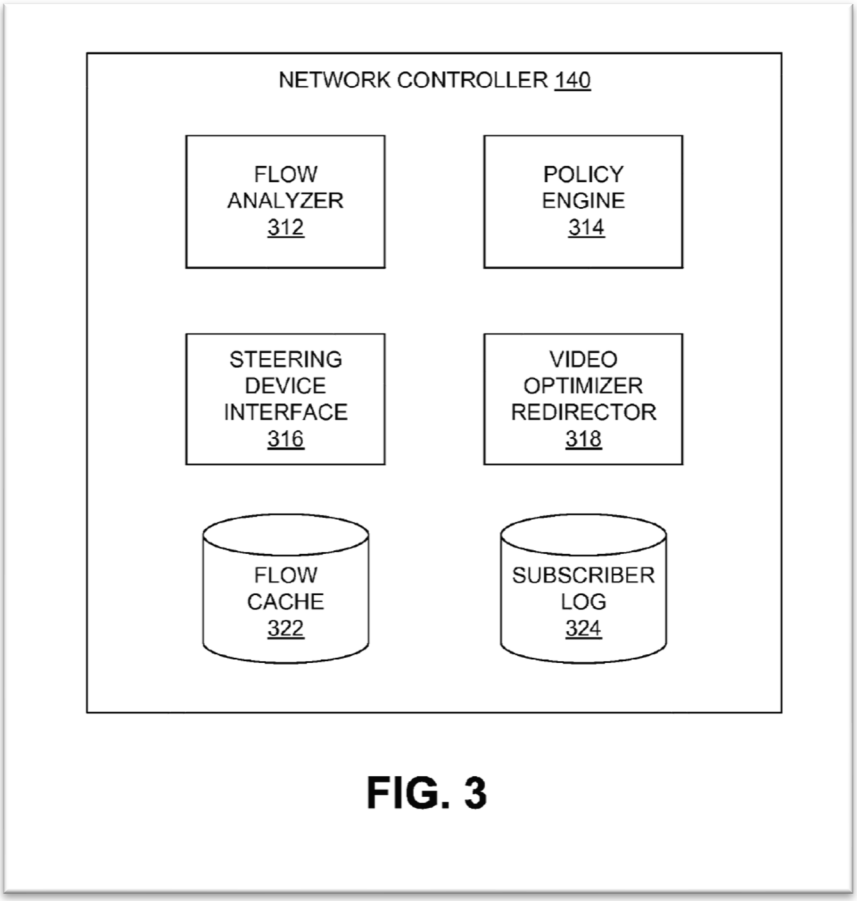


Figure 3 of Swenson

Swenson also explains that "[t]he flow cache entry keeps track of the flow and its associated bandwidth." EX1007, ¶[0061]; *see id.*, ¶[0040], ¶[0046], ¶[0062]. In discussing Figure 7, Swenson states that the flow cache entry stores at least a portion of the received packet, including a least the IP address of the received packet. *Id.*, ¶[0073] (discussing Figure 7); *see id.*, ¶[0074]-¶[0085]. Further, Swenson states that a flow cache entry can also store the MAC address or TCP source port associated with a received packet. *Id.*, ¶[0084]. A POSA would have understood from these disclosures that Swenson teaches that the controller stores at least a portion of a packet in a flow cache memory. EX1004, ¶¶134-135.

G. Claim 7

[7] The method according to claim 5, further comprising responsive to the packet satisfying the criterion and to the instruction, sending a portion of the packet, by the network node, to the controller.

Claim 7 is obvious over Swenson. EX1001, 11:34-37.

As discussed above for Element [1.0] and Claim 5, Swenson teaches that the network node sends a portion of a packet to the controller in response to the packet satisfying a criterion. Swenson discloses that, when its system detects a HTTP packet flow matching a particular signature, "the steering device 130 **forwards the HTTP request and a portion of the HTTP response to the network controller 140** over the ICAP client interface 404." EX1007, ¶[0059] (emphasis added]; *see id.*, ¶[0060], ¶[0065], Figures 1, 4A-4B. A POSA would have understood that the steering device

130 corresponds to at least part of the claimed network node, and the signature that is compared to the HTTP packet flow corresponds to the claimed criterion. EX1004, ¶¶136-137. Further, a POSA would have understood that, responsive to the HTTP packet flow satisfying that criterion, the steering device applies an instruction pursuant to which a portion of a HTTP request or a portion of a HTTP response is sent to the network controller. EX1004, ¶¶137-138.

H. Claim 8

[8] The method according to claim 7, wherein the portion of the packet consists of multiple consecutive bytes, and wherein the instruction comprises identification of the consecutive bytes in the packet.

Claim 8 is obvious over Swenson in view of the knowledge of a POSA. EX1001, 11:38-41. As discussed above for Claim 7, Swenson discloses that a portion of the HTTP request or the HTTP response is sent to the controller. EX1007, ¶[0059]-¶[0060], Figures 1, 4A-4B. It would have been obvious to a POSA that the portion of the HTTP request or the HTTP response sent to the controller consists of multiple consecutive bytes in the packet that were identified by an instruction implemented in steering device 130 of Swenson. EX1004, ¶¶139-141.

I. Claim 9

[9] The method according to claim 5, further comprising responsive to receiving the packet, analyzing the packet, by the controller.

Claim 9 is obvious in view of Swenson. EX1001, 11:42-44.

As discussed above for Element [1.0], Claim 5 and Claim 7, Swenson discloses that packets are sent to the network controller. Swenson also explains that, after receiving the packet, the flow analyzer 312 of the network controller analyzes the packet through DPI:

> After receiving the request and the portion of response at the ICAP server interface 406, **the flow analyzer 312 of the network controller 140 performs a deep flow inspection** to determine if the flow is worth bandwidth monitoring and/or user detection.

EX1007, ¶[0059] (emphasis added); *see id.*, ¶[0060] (stating that the "controller 140 ingests the network flow for inspection"), Figures 1, 4A-4B; EX1004, ¶¶142-143. Further, Swenson discusses "an example event trace of [Swenson's] 'continue' working mode" in which the steering device 130 "sends an ICAP request message 516 comprising [a] HTTP GET request header and a portion of the [HTTP] response payload to the network controller 140, which inspects the message to determine whether to monitor the flow or optimize the video." *Id.*, ¶[0065], Figure 5. A POSA would have understood the inspection of the "HTTP GET request header and a portion of the [HTTP] response payload" to be analyzing the packet upon the controller receiving the packet. EX1004, ¶143. This is similar to Lin, which explains that an analysis is performed on its packets. EX1005, 1:58-2:4, 3:25-33, 4:8-31, 4:61-65, 5:26-36, 5:45-51, 6:1-12, 6:50-57, 7:59-61, 8:12-16, Figures 2, 5-9; EX1004, ¶¶143-144.

J. Claim 12

[12] The method according to claim 9, wherein the analyzing comprises applying security or data analytic application.

Claim 12 is obvious over Lin. EX1001, 11:55-57. Lin states that the analysis performed on the packets includes a security processing function. EX1005, 1:58-2:4, 3:25-33, 4:8-31, 4:61-65, 5:26-36, 5:45-51, 6:1-12, 6:50-57, 7:59-61, 8:12-16, Figures 2, 5-9; EX1004, ¶¶145-148.

K. Claim 13

[13] The method according to claim 9, wherein the analyzing comprises applying security application that comprises firewall or intrusion detection functionality.

Claim 13 is obvious over Lin. EX1001, 11:58-60. Lin discloses that its system analyzes packets by applying a security application that comprises either a firewall or intrusion detection functionality, or both. EX1005, 5:45-50 (referring specifically to "compliance with firewall rules" and "network intrusion detection"**)**, 3:11-12, 6:50-54; EX1004, ¶¶149-151.

L. Claim 14

[14] The method according to claim 9, wherein the analyzing comprises performing Deep Packet Inspection (DPI) or using a DPI engine on the packet.

Claim 14 is obvious over both Lin and Swenson. EX1001, 11:61-63. Lin states that its analysis includes "deep packet inspection (DPI)." EX1005, 3:11-12; EX1004, ¶¶152-155. As discussed above for Element [1.0], Swenson also discloses

that the flow analyzer in network controller 140 performs DPI. EX1007, ¶[0059]-¶[0060], ¶[0065].

M. Claim 15

[15] The method according to claim 9, wherein the packet comprises distinct header and payload fields, and wherein the analyzing comprises checking part of, or whole of, the payload field.

Claim 15 is obvious in view of both Lin and Swenson. EX1001, 11:64-67.

A POSA would have known that a packet includes distinct header and payload fields. EX1008, Section 2; EX1004, ¶¶156-157.

Further, Swenson disloses "an example event trace of [Swenson's] 'continue' working mode" in which the steering device 130 "sends an ICAP request message 516 comprising [a] HTTP GET request header and a portion of the [HTTP] response payload to the network controller 140, which inspects the message to determine whether to monitor the flow or optimize the video." EX1007, ¶[0065], Figure 5; EX1004, ¶158. A POSA would have understood the analysis of the "response payload" in Swenson to be checking part of the payload field of a packet by the network controller. Further, a POSA would have understood the references to DPI in both Swenson and Lin to involve checking part or the whole of the payload field. EX1004, ¶¶159-160; EX1005, 3:11-12, 5:45-50; EX1007, ¶[0059]-¶[0060].

N. Claim 16

[16] The method according to claim 1, wherein the packet comprises distinct header and payload fields, the header comprises one or more flag bits, and

wherein the packet-applicable criterion is that one or more of the flag bits is set.

Claim 16 is obvious over Lin in view of the knowledge of a POSA. EX1001, 12:1-4. A POSA would have known that a packet includes distinct header and payload fields. EX1008, Section 2; EX1004, ¶¶161-162. Further, it would have been well-known to a POSA that the header field includes one or more flag bits. EX1004, ¶¶162-163; EX1008, Section 2. In addition, it would have been obvious to a POSA that the flow rules of Lin could have used a determination of whether or not a flag bit was set as one of Lin's "conditions" that serve as packet-applicable criterion. EX1004, ¶¶162-163; EX1005, 4:23-31; *see id.*, 5:8-12, 7:24-8:18.

O. Claim 17

[17] The method according to claim 16, wherein the packet is an Transmission Control Protocol (TCP) packet, and wherein the one or more flag bits comprises comprise a SYN flag bit, an ACK flag bit, a FIN flag bit, a RST flag bit, or any combination thereof.[3]

Claim 17 is obvious over Lin in view of the knowledge of a POSA. EX1001, 12:5-9. Lin discloses embodiments in which the packet is a TCP packet. EX1005, 7:24-8:18, Tables 2-3. Further, it would have been well-known to a POSA that flag bits such as a SYN flag bit, an ACK flag bit, a FIN flag bit, and a RST flag bit were

[3] Claim 17 uses the language "comprises comprise." This may be an error. For this Petition, Petitioner interprets that claim language to mean "comprises."

commonly used as part of the "handshake" process to initiate transmission of TCP packets (in the case of SYN and ACK flag bits) and to reset or terminate a TCP connection (in the case of RST and FIN flags, respectively). EX1004, ¶¶164-165. As such, it would have been obvious to a POSA that the flow rules of Lin could have used a determination of whether or not one or more of these flag bits was set as one of Lin's "conditions" that serve as packet-applicable criterion. EX1004, ¶¶165-166; EX1005, 4:23-31; *see id.*, 5:8-12, 7:24-8:18.

P. Claim 18

[18] The method according to claim 1, wherein the packet comprises distinct header and payload fields, the header comprises at least the first and second entities addresses in the packet network, and wherein the packet-applicable criterion is that the first entity address, the second entity address, or both match a predetermined address or addresses.

Claim 18 is obvious over Lin in view of the knowledge of a POSA. EX1001, 12:10-16. A POSA would have known that a packet includes distinct header and payload fields. EX1008, Section 2; EX1004, ¶¶167-168. Moreover, it would have been well-known that the header would include the IP address of the source of the packet and the IP address of the destination of the packet. EX1004, ¶168. Further, as disclosed above for Element [1.1], Lin provides an example where the packet-applicable criterion is that the address of the first entity ("MAC src" or "IPSrc"), the address of the second entity ("MAC dst" or "IP dist"), or both match a predetermined

address or addresses to identify a particular packet. EX1005, 5:8-25; *see id.*, 5:26-36, 6:1-12, 6:40-54, Table 1; EX1004, ¶¶169-170.

Q. Claim 19

[19] The method according to claim 18, wherein the addresses are Internet Protocol (IP) addresses.

Claim 19 is obvious over both Lin for the reasons as those discussed above for Claim 18. EX1001, 12:17-18; EX1005, 5:8-36, 6:1-12, 6:40-54, Table 1; EX1004, ¶¶171.

R. Claim 20

[20] The method according to claim 1, wherein the packet is an Transmission Control Protocol (TCP) packet that comprises source and destination TCP ports, a TCP sequence number, and a TCP sequence mask fields, and wherein the packet-applicable criterion is that the source TCP port, the destination TCP port, the TCP sequence number, the TCP sequence mask, or any combination thereof, matches a predetermined value or values.

Claim 20 is obvious over Lin in view of the knowledge of a POSA. EX1001, 12:20-27.

Lin provides examples of TCP packets that are transported over a TCP network when discussing Tables 2 and 3 of that patent. EX1005, 7:24-8:18; *see id.*, 10:48-50; EX1004, ¶¶172-173. Further, Lin discloses tracking HTTP packets whose source and destination TCP ports are port 80. EX1005, 7:24-8:18. Lin does not specifically disclose that a TCP packet has a TCP sequence number or a TCP

sequence mask field, but this would have been obvious to a POSA because those fields are always present in a TCP packet. EX1004, ¶173.

In addition, Lin explains that the packet applicable criterion used to determine whether to take an action on a packet can include either the source TCP port or destination TCP port. EX1004, ¶¶174-175. Lin's discussion of Table 3 (shown below) teaches that the devices check for a specific packet-applicable criterion – that port 80 is the source or destination port of the TCP packets, which indicates that they are HTTP packets. EX1005, 7:64-8:18. Lin states that, if a packet satisfies this criterion (i.e., the packet indicates that port 80 is the source or destination port), then the SDN switch (which corresponds to the claimed network node) sends the packet over the packet network to the security device (which corresponds to the entity that is included in the instruction and is other than the second entity). *Id.*, 8:10-18.

TABLE 3

IN_PORT	. . .	IP src	TCP src port	TCP dst port	. . .	Action	Count
Ingress_port_ID	*	*	*	80	*	Redirect port	10
Redirect_port_ID	*	*	80	*	*	Ingress port	10
Ingress_port_ID	*	*	*	*	*	Egress port	130
Egress_port_ID	*	*	*	*	*	Ingress port	130

Table 3 of Lin (Annotated)

Id., 8:10-18 (emphasis added), Table 3.

S. Claim 21

[21] The method according to claim 1, wherein the packet network comprises a Wide Area Network (WAN), Local Area Network (LAN), the Internet, Metropolitan Area Network (MAN), Internet Service Provider (ISP) backbone, datacenter network, or inter-datacenter network.

Claim 21 is obvious over Lin in view of the knowledge of a POSA. EX1001, 12:28-32. As discussed above for Claims 18-19, Lin discloses the use of IP and MAC addresses to identify particular packets. EX1005, 5:8-36, 6:1-12, 6:40-54, Table 1. A POSA would have understood that, if IP and MAC addresses were being used in Lin, then the packet network would have included one of the Internet, a Wide Area Network ("WAN"), a Local Area Network ("LAN"), a Metropolitan Area Network ("MAN"), Internet Service Provider ("ISP") backbone, datacenter network, or inter-datacenter network. EX1004, ¶¶176-178.

T. Claim 22

[22] The method according to claim 1, wherein the first entity is a server device and the second entity is a client device, or wherein the first entity is a client device and the second entity is a server device.

Claim 22 is obvious over Swenson. EX1001, 12:33-36. Swenson teaches "user devices 110" on one end of its system and an "origin server 160" on the other end of its system. EX1007, ¶[0023]; EX1004, ¶¶179-180. Swenson discloses that the user devices "are computing devices with network capabilities," such as "laptops, notebooks, tablets, smart telephones, or personal digital assistants (PDAs)." EX1007, ¶[0025]. A POSA would have understood these to be client

devices. EX1004, ¶¶180-181. Given that Paragraph 58 of Swenson discloses two way communication between the user devices and the origin server, Swenson identifies user devices 110 (i.e., client devices) as the claimed first entity, and origin server 160 (i.e., a server) as the claimed second entity, or vice versa. EX1007, ¶[0058]; EX1004, ¶180.

U. Claim 23

[23] The method according to claim 22, wherein the server device comprises a web server, and wherein the client device comprises a smartphone, a tablet computer, a personal computer, a laptop computer, or a wearable computing device.

Claim 23 is obvious over Swenson. EX1001, 12:37-41. As discussed above for Claim 22, Swenson teaches "user devices 110" on one end of its system and an "origin server 160" on the other end of its system. EX1007, ¶[0023]. A POSA would have understood that origin server 160 comprises a web server because origin server 160 is able to process HTTP requests, which are requests for web content. EX1007, ¶[0058]; EX1004, ¶¶182-184. Further, Swenson discloses that the user devices (i.e., client devices) can include "laptops, notebooks, tablets, smart telephones, or personal digital assistants (PDAs)." EX1007, ¶[0025].

V. Claim 24

[24] The method according to claim 22, wherein the communication between the network node and the controller is based on, or uses, a standard protocol.

Claim 24 is obvious in view of Lin. EX1001, 12:42-44. Lin discloses the use of the OpenFlow protocol for the communication between the SDN switch (i.e., the network node) and the SDN controller. EX1005, 1:18-43, 2:12-13, 3:42-46, 3:65-4:6, 6:35-39; EX1004, ¶¶185-187. Figure 1 of Lin (shown below) describes the use of OpenFlow protocol as "prior art" to Lin, which indicates that OpenFlow protocol was known as a standard protocol well before the priority date for the '111 Patent.

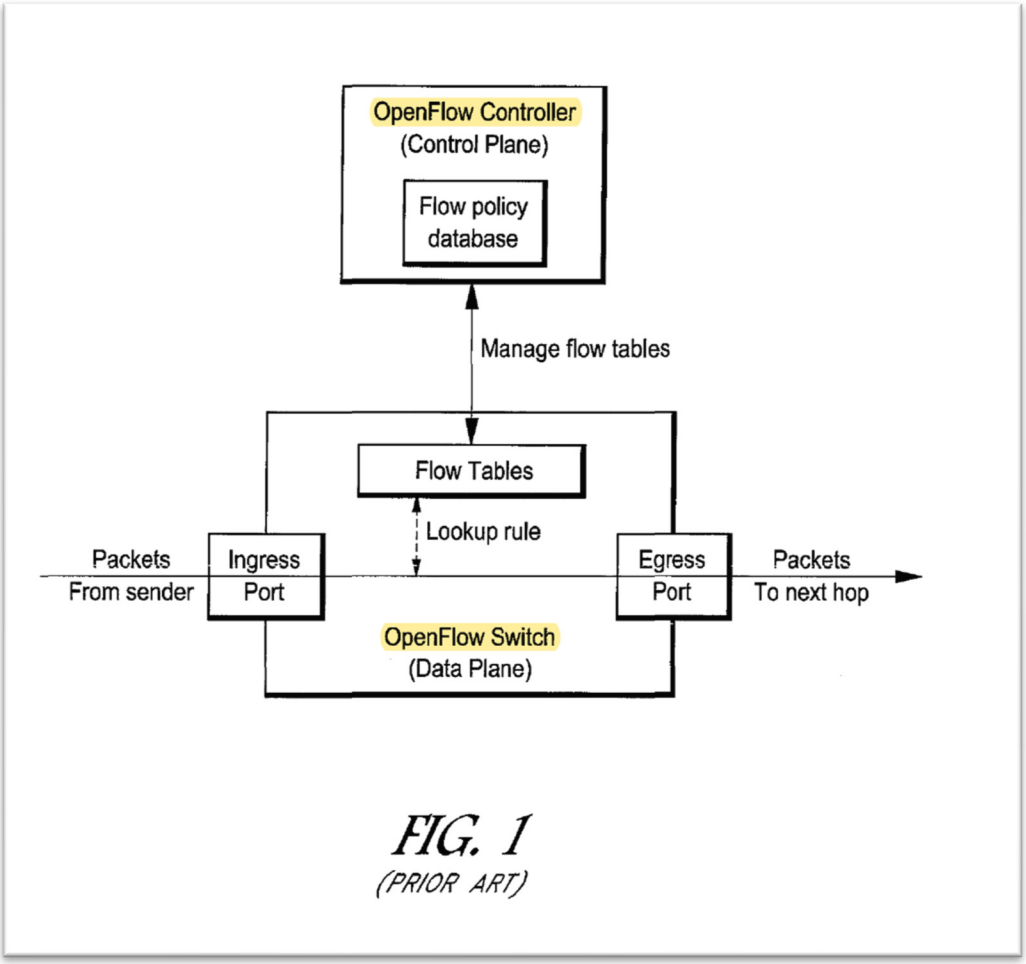


Figure 1 of Lin (Annotated)

EX1005, Figure 1.

W. Claim 27

[27] The method according to claim 1, wherein the network node comprises a router, a switch, or a bridge.

Claim 27 is obvious in view of Lin. EX1001, 12:51-52. As discussed above for Element [1.0], Lin discloses a SDN switch that operates as the claimed network node. EX1005, 1:58-2:4, 4:33-67, 6:13-23, 6:57-63; EX1004, ¶¶188-190.

X. Claim 28

[28] The method according to claim 1, wherein the packet network is an Internet Protocol (IP) network, and the packet is an IP packet.

Claim 28 is obvious in view of Lin. EX1001, 12:53-55. As discussed above for Claims 18-19, Lin discusses the use of IP addresses to identify packets being transmitted across the packet network. EX1005, 5:8-36, 6:1-12, 6:40-54, Table 1. A POSA would have understood from the use of the IP addresses that packet network in Lin is an Internet Protocol (IP) network and the packet is an IP packet. EX1004, ¶¶191-193.

Y. Claim 29

[29] The method according to claim 28, wherein the packet network is an Transmission Control Protocol (TCP) network, and the packet is an TCP packet.

Claim 29 is obvious in view of Lin. EX1001, 12:56-58. Lin provides specific examples of TCP packets that are transported over a TCP network when discussing

Tables 2 and 3 of that patent. EX1005, 7:24-8:18; *see id.*, 10:48-50; EX1004, ¶¶194-196.

Z. Claim 30

[30.0] The method according to claim 1, further comprising:

receiving, by the network node from the first entity over the packet network, one or more additional packets; checking, by the network node, if any one of the one or more additional packets satisfies the criterion;

Element [30.0] is obvious in view of Lin. EX1001, 12:59-63. Lin teaches that its method can be applied to "any packet received by the SDN switch." EX1005, 6:40-63. As such, Lin would check if one or more additional packets after the first packet satisfies the packet-applicable criterion for the reasons discussed above for Elements [1.4]-[1.6]. EX1004, ¶¶197-199.

[30.1] responsive to an additional packet not satisfying the criterion, sending, by the network node over the packet network, the additional packet to the second entity; and responsive to the additional packet satisfying the criterion, sending the additional packet, by the network node over the packet network, in response to the instruction.

Element [30.1] is obvious in view of the combination of Lin and Shieh. EX1001, 12:64-13:3. Lin discloses that the SDN switch forwards additional packets over the packet network in response to whether not the packets satisfy the packet-applicable criterion for the reasons discussed above for Elements [1.5]-[1.6]. EX1004, ¶¶200-203.

AA. Claim 31

[31] The method according to claim 1, wherein the packet network is a Software Defined Network (SDN), the packet is routed as part of a data plane and the network node communication with the controller serves as a control plane.

Claim 31 is obvious in view of Lin. EX1001, 13:4-7. Lin explains that its packet network is a SDN network that uses an SDN controller and an SDN switch. EX1005, Abstract, 1:7-9, 1:58-2:4, 3:53-4:67, Figures 2, 6-8. 2, 6-8. Lin also discloses that it uses an Openflow protocol in which a packet is routed as part of the data plane through its SDN network. *Id.*, 1:11-54, 2:49-50, 3:42-52, Figure 1; EX1004, ¶¶204-205. Further, Lin explains that, when using the Openflow protocol, the communication between the SDN switch and the SDN controller is through a control plane. EX1005, 1:11-54, 2:49-50, 3:42-52, 3:65-4:5, Figure 1; EX1004, ¶¶205-206.

X. GROUND 2: CLAIMS 1, 5-9, 12-24 AND 27-30 ARE UNPATENTABLE AS OBVIOUS OVER SHIEH IN VIEW OF SWENSON.

The combination of Shieh and Swenson, along with the knowledge of a POSA, renders Claims 1, 5-9, 12-24 and 27-30 obvious. EX1004, ¶¶207-318.

A. Claim 1

Element [1.0]

To the extent Element [1.0] is limiting, it is obvious in view of Shieh. EX1001, 10:52-55.

Shieh discloses a method for use with a packet network. EX1004, ¶¶208-209. Shieh's Abstract states that it relates to "[a] network system" that operates on a "packet from a source node destined to a destination node." EX1005, Abstract. The specification of Shieh further explains that its system performs operations on packets passing through a network. EX1006, ¶[0002], ¶[0017], ¶[0018], ¶[0021], ¶[0023]. Further, Figure 7 of Shieh, reproduced below, shows the method disclosed in Shieh, which involves a "network access device" performing operations on packets passing through that network access device. *Id.*, ¶[0049], Figure 7.

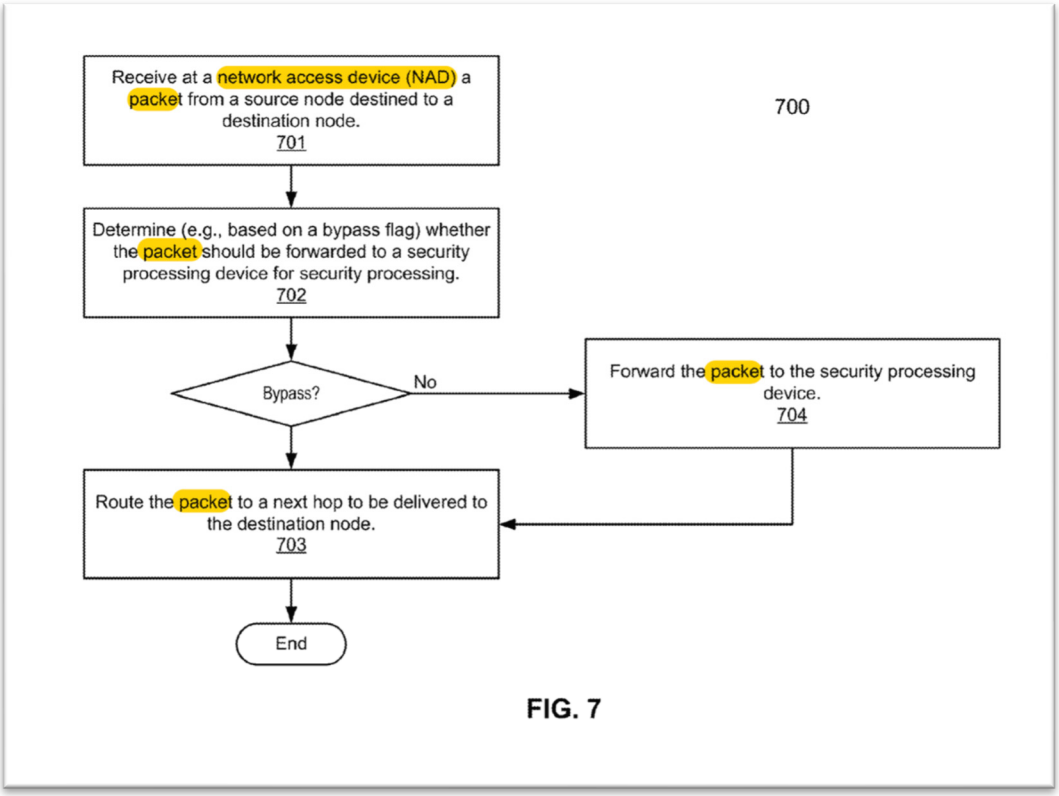


Figure 7 of Shieh (Annotated)

Further, Shieh discusses a network node for transporting packets between first and second entities. EX1004, ¶210. Shieh discloses "network access devices 204A-204C" (**each of which corresponds to the claimed network node**) for transporting a packet from a "source computer" (**which corresponds to the claimed first entity**) to a "destination, which may be another host, a multicast group or a broadcast domain" (**any of which corresponds to the claimed second entity**). EX1006, ¶[0027]; ¶[0037], ¶[0049], Figures 2B, 7. The way in which packets pass from a source node through network access devices 204A-204C to a destination node can be seen below in Figure 2A of Shieh:

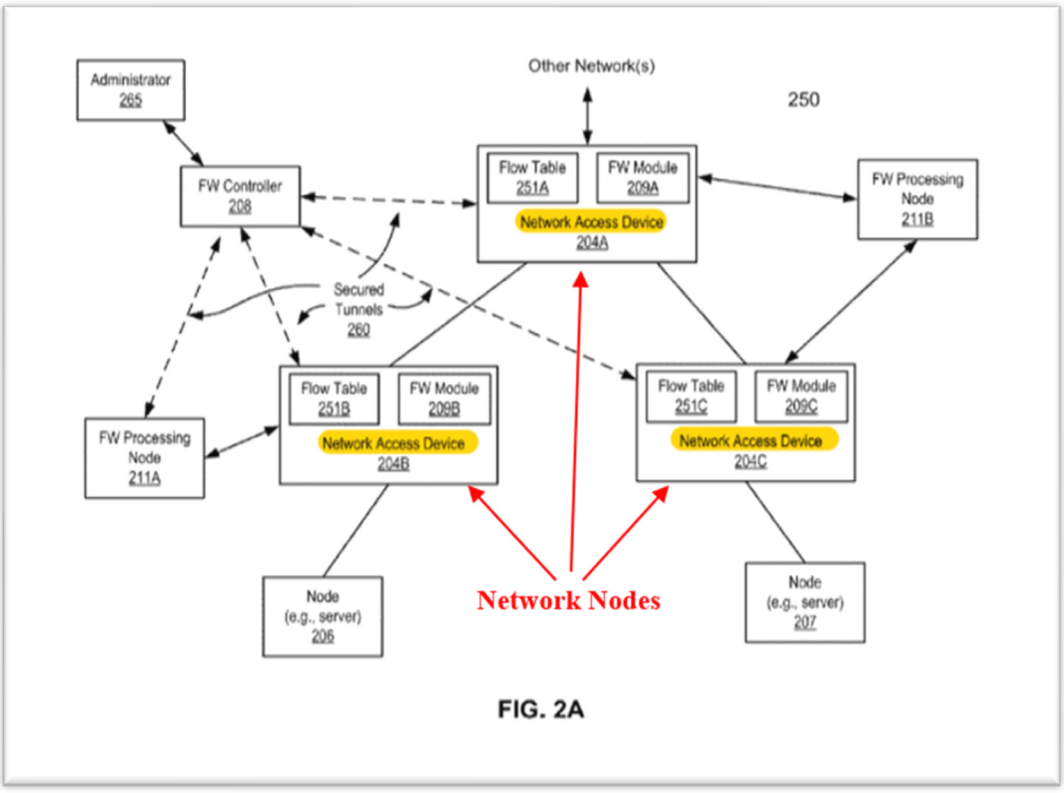


Figure 2A of Shieh (Annotated)

Id., Figure 2A.

Moreover, Shieh explains that the network nodes (e.g., network access devices 204A-204C) are under the control of a controller (**which corresponds to the claimed controller**) that is external to the network node. EX1004, ¶211. Shieh explains that network access devices 204A-204C are "configured" by (i.e., under the control of) "a controller or a management entity." EX1006, ¶[0018]; *see id.*, ¶[0021] (noting that firewall controller 208 may be "external to network access device 204"), ¶[0025], ¶[0029], ¶[0042]. Figure 2A shows that the controller 208 is external to the network node:

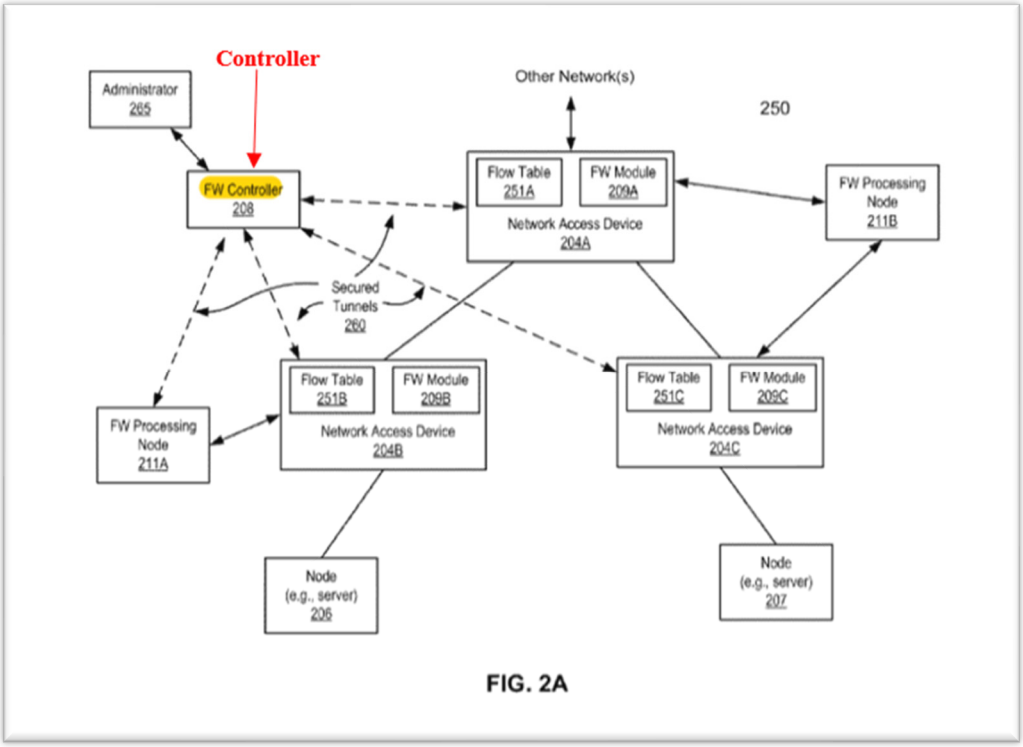


Figure 2A of Shieh (Annotated)

Id., Figure 2A, ¶[0029]; *see id.*, ¶[0025]-¶[0026].

In addition, as addressed in the above claim construction section, the claimed controller is an entity configured to perform DPI on packets.[4] EX1004, ¶212. The combination of Shieh and Swenson renders it obvious that the controller is configured to perform DPI. EX1004, ¶¶212-214. Shieh discloses that its security processing function incudes DPI. EX1006, ¶[0021]. Moreover, as discussed above in Ground 1, Swenson discloses a controller configured to perform DPI. *See* EX1007, ¶[0059] ("the flow analyzer 312 of the network controller 140 performs a deep flow inspection …"), ¶[0060], (stating that the "controller 140 ingests the network flow for inspection"), ¶[0065] (discussing inspection of HTTP response payload by network controller 140), Figures 1, 4A-4B; EX1004, ¶¶213-214.

Given the similarities between the disclosures in Lin and Shieh, a POSA would have been motivated to combine the teachings of Shieh and Swenson for similar reasons as those disclosed above in Ground 1 for the combination of Lin and Swenson. EX1004, ¶¶82-86, 215-220. For instance, Shieh and Swenson relate to routing traffic through a computer network and use a central controller to provide instructions and packet-applicable criterion to network nodes to determine which

[4] To the extent that the PTAB does not agree with this construction, Shieh still discloses Element [1.0] for the reasons discussed above.

packets should be redirected and which packets can be sent directly to a destination node. EX1006, ¶[0017]-¶[0018], ¶[0024]-¶[0025], ¶[0029]-¶[0030], Figures 1-3, 7; EX1007, ¶[0023]-¶[0032], ¶[0038]-¶[0043], ¶[0057]-¶[0061], Figures 1-4A; EX1004, ¶215.

Further, a POSA would have understood both Shieh and Swenson to disclose redirecting packets for analysis as part of a security function. EX1004, ¶¶216-217. Shieh states that a "packet is transmitted to a security device associated with the network access device to allow the security device to perform content inspection." EX1006, Abstract; *see id.*, ¶[0042], ¶[0049]. Similarly, a POSA would have understood that the bandwidth monitoring in Swenson can be used as a security application that monitors for DOS attacks that occupy significant bandwidth in a network. EX1007, ¶[0059]-¶[0060]; *see id.*, ¶[0039] (explaining that the controller can incorporate "security functions"; EX1004, ¶216. Further, the architecture of Swenson is functionally similar to the architecture of Shieh, as can be seen in the below two figures showing a central controller directly connected to a network node (the Network Access Devices in Shieh and the steering device in Swenson) through which packet flows pass:

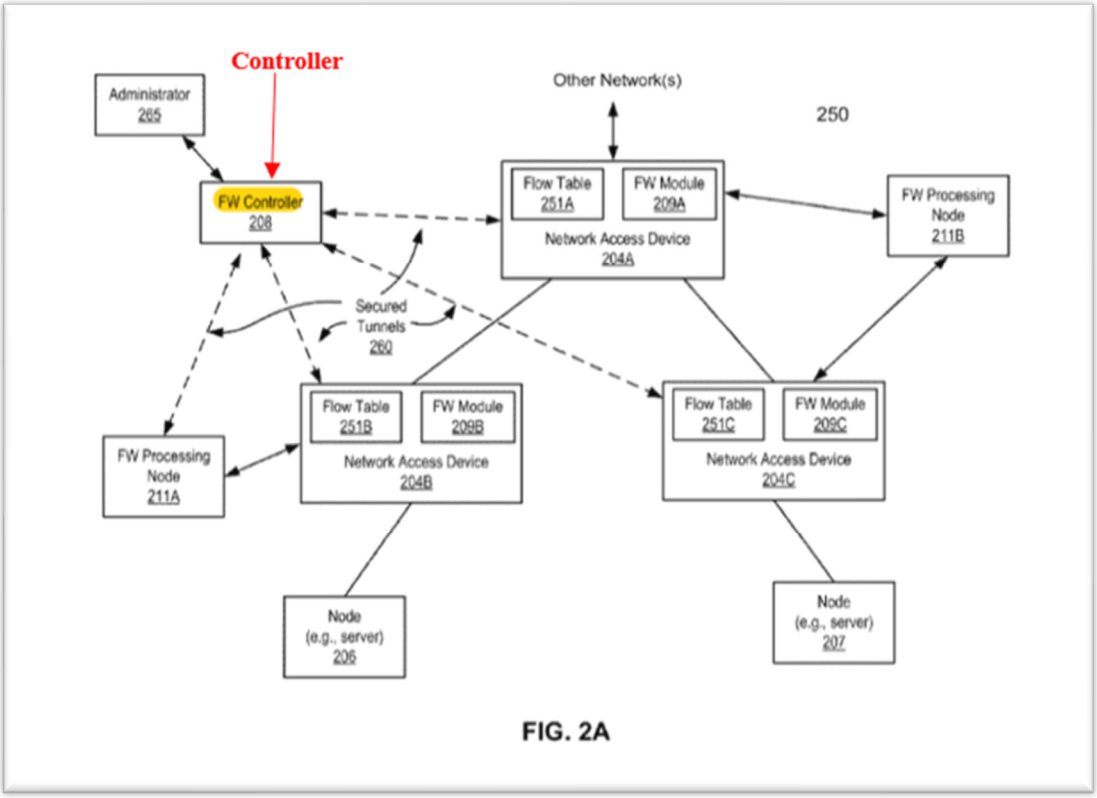


Figure 2A of Shieh (Annotated)

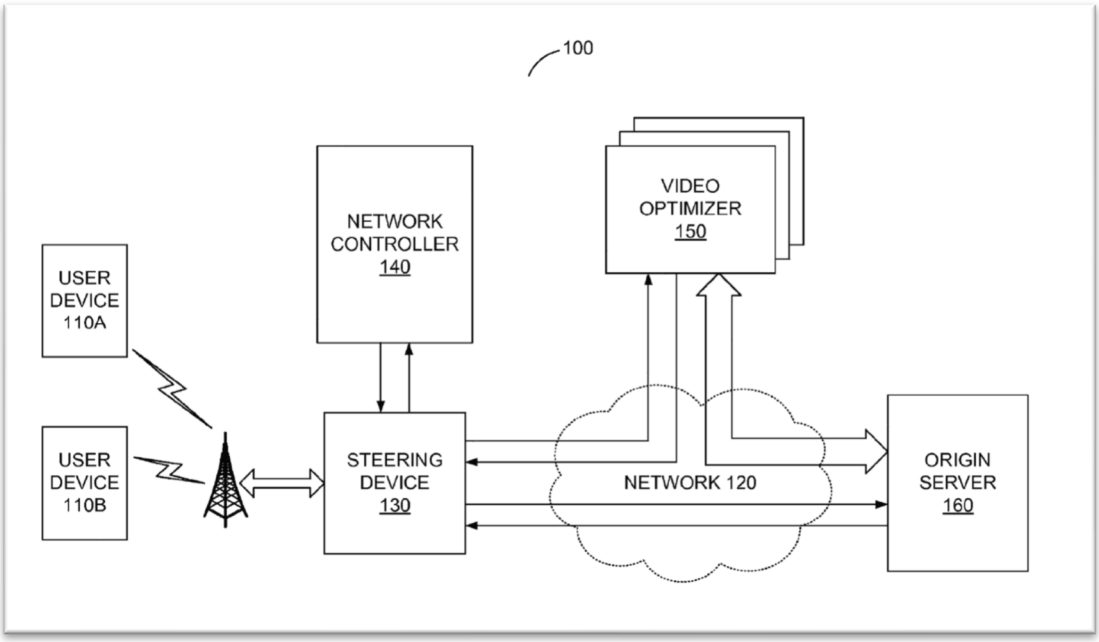


Figure 1 of Swenson

EX1004, ¶216. In addition, Shieh and Swenson each teach that only certain packets may need to be redirected for further processing before being sent to the destination node. EX1006, ¶[0017]-¶[0018], EX1007, ¶[0023]-¶[0032].

Further, a POSA would have found it obvious to implement the security processing module in Shieh as part of the controller in light of the disclosures in Swenson to send packets to the controller for DPI. EX1007, ¶[0059]-¶[0060], Figures 1, 4A-4B; EX1004, ¶218. The controller provides a central location for packets from different nodes to undergo inspection by a security component that applies the same security algorithms to each analyzed packet. EX1004, ¶218. Further, a POSA would have known that security algorithms in a security component are often updated via software updates as new threats are identified. *Id.* A POSA would have known that an efficient way to keep the security component up-to-date was to have a central security component that is part of the central processor. *Id.*

Moreover, a POSA would have had a reasonable expectation of success in modifying Shieh to implement the disclosure of Swenson given that they have a functionally similar architecture. EX1006, ¶[0017]-¶[0018], ¶[0024]-¶[0025], ¶[0029]-¶[0030], Figures 1-3, 7; EX1007, ¶[0023]-¶[0032], ¶[0038]-¶[0043], ¶[0057]-¶[0061], Figures 1-4A; EX1004, ¶219. If Swenson could be constructed to have the network nodes send packets to the central controller for inspection, a POSA

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 294 of 496

would have understood that the same arrangement could be implemented in Shieh. EX1004, ¶219.

Element [1.1]

Element [1.1] is obvious in view of Shieh. EX1001, 10:56-58.

Shieh discloses that Controller 208 controls network access devices 204A-204C (i.e., the claimed network nodes) in a packet network for the reasons discussed above for Element [1.0]. EX1006, ¶[0018], ¶[0025]-¶[0026], ¶[0029]; EX1004, ¶¶221-222.

Further, Shieh identifies an instruction sent by the controller to the network node over the packet network that determines whether or not a packet requires inspection. EX1004, ¶223. Shieh states that "a persistent connection" exists over the packet network to allow for communication between the controller and the network access devices 204A-204C. EX1006, ¶[0025]; *see id.*, ¶[0018], ¶[0028]-¶[0029]. Further, Shieh discloses that a "command" (**which corresponds to the claimed instruction**) is sent from an administrator via the controller to the network access devices 204A-204C "to set up a set of filtering rules concerning whether and/or what types of packets should be forwarded to a security device." *Id.*, ¶[0017]-¶[0018]; *see id.*, ¶[0023], ¶[0028]-¶[0029].

Shieh also discloses separate packet-applicable criterion to determine which packets should be forwarded to the security device pursuant to the filtering rules.

EX1004, ¶224. For example, Shieh provides criterion specific to a particular packet being checked, such as the identification of "TCP FIN or TCP RST packets." *Id.*, ¶[0036]; *see id.*, ¶[0035], ¶[0049]. **These correspond to the claimed packet-applicable criterion.** Shieh does not state whether these criterion are sent from a controller to the network access devices. However, a POSA would have understood that the packet-applicable criterion would be sent from a controller to one or more network access devices. EX1004, ¶¶224-225. This allows the same packet-applicable criterion to be applied at multiple network access devices (i.e., network nodes). *Id.*

Element [1.2]

Element [1.2] recites that the instruction and packet-applicable criterion that are sent by the controller to the network node in Element [1.1] are received by the network node. EX1001, 10:59-60. Shieh therefore discloses Element [1.2] for substantially the same reasons that it discloses Element [1.1]. EX1004, ¶¶226-232.

Element [1.3]

Element [1.3] is obvious in view of Shieh. EX1001, 10:60-62.

Claim 1 of Shieh is a method claim that discloses that a network access device (i.e., a network node) receives a packet from a source node (i.e, a first entity) that is addressed to a destination node (i.e., a second entity). EX1006, Claim 1; EX1004, ¶¶233-235; *see* EX1006, Abstract, ¶[0020], ¶[0027], ¶[0037], ¶[0049]. Shieh

explains that the packets may be addressed with, among other things, "Source and Destination IP address," which are used to address the source computer (i.e., the claimed first entity) and the destination (i.e., the claimed second entity). *Id.*, ¶[0027]. In addition, Figure 1 of Shieh (reproduced below) shows a system in which "Network Access Device(s) 204" receive packets from client 201 or client 202 that are addressed to server 206 or server 207:

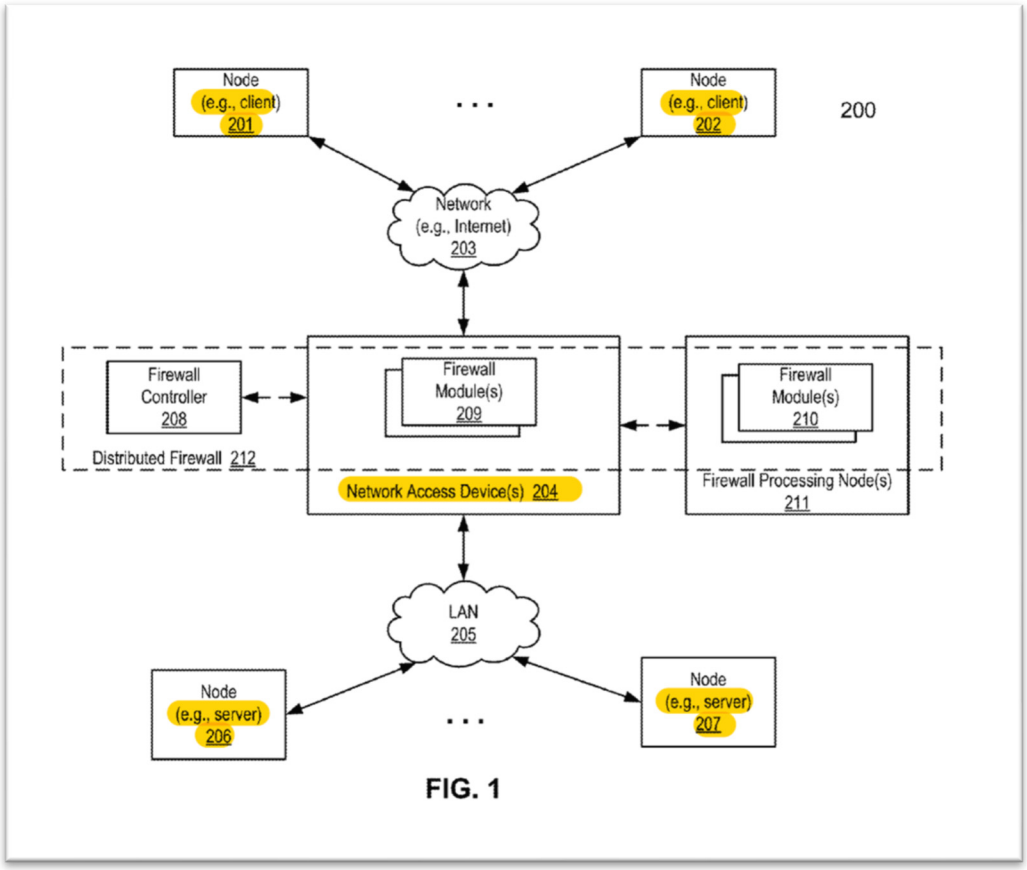


Figure 1 of Shieh (Annotated)

Id., Figure 1; *see id.*, Figures 2A-2B, 7.

Element [1.4]

Element [1.4] is obvious in view of Shieh. EX1001, 10:63-64.

As discussed above for Element [1.1], Shieh discusses the implementation of packet-applicable criterion, such as the identification of "TCP FIN or TCP RST packets," which are checked by the network access device to determine whether a packet satisfies that criterion. EX1006, ¶[0036]; *see id.*, ¶[0035], ¶[0037], ¶[0049], Figure 7; EX1004, ¶¶236-238.

Element [1.5]

Element [1.5] is obvious in view of Shieh. EX1001, 10:65-67.

As discussed above for Elements [1.1] and [1.4], Shieh discloses criterion specific to a particular packet being checked by the network access device. EX1006, ¶[0018], ¶[0023], ¶[0028]-¶[0029], ¶[0035], ¶[0049], Figure 7. Shieh checks whether a packet is a TCP FIN or TCP RST packet, and, responsive to the packet not satisfying this criterion, sends the packet to the destination node (i.e., the second entity). *Id.*, ¶[0036]; *see id.*, ¶[0037] (discussing whether packet should be sent to security processing device "based on a bypass flag"); EX1004, ¶¶239-242.

Element [1.6]

Element [1.6] is obvious in view of Shieh. EX1001, 11:1-4.

As discussed above for Elements [1.1], [1.4] and [1.5], Shieh teaches that its network access device utilizes packet-applicable criterion to determine whether a

packet should be forwarded to a security device. EX1006, EX1006, ¶[0018], ¶[0023], ¶[0028]-¶[0029], ¶[0035], ¶[0049], Figure 7; EX1004, ¶¶243-244. Shieh checks whether a packet is a TCP FIN or TCP RST packet, and, responsive to the packet satisfying this criterion, sends the packet to the security device (i.e., the entity that is included in the instruction and is other than the second entity). EX1006, ¶[0036]; *see id.*, ¶[0037] (discussing whether packet should be sent to security processing device "based on a bypass flag"); EX1004, ¶¶244-245.

B. Claims 5-9

Claims 5 and 7 are obvious over the combination of Shieh (which discloses the instruction and packet-applicable criterion for the reasons discussed above in Elements [1.1]-[1.2] and [1.4]-[1.6]) and Swenson (which discloses sending the packet or a portion thereof to the controller for the reasons discussed in Ground 1). EX1004, ¶¶246-249, 253-255.

Claims 6 and 8-9 are obvious over Swenson, in view of the knowledge of a POSA, for the reasons discussed above in Ground 1. EX1004, ¶¶250-252, 256-262.

C. Claim 12

Claim 12 is obvious over Shieh. EX1001, 11:55-57. Shieh repeatedly states that the analysis performed on the packets includes a security processing function. EX1006, ¶[0002], ¶[0017]-¶[0019], ¶[0021], ¶[0023], ¶[0029]-¶[0031], ¶[0035]-

¶[0037], ¶[0042], ¶[0049], Figures 1, 2B, 3, 7. A POSA would have understood from this that the analysis includes the security processing function. EX1004, ¶¶263-266.

D. Claim 13

Claim 13 is obvious over Shieh. EX1001, 11:58-60. Shieh discloses that its system analyzes packets by applying a security application that comprises a firewall functionality. EX1006, ¶[0021], ¶[0023], ¶[0027]-¶[0028], ¶[0038], ¶[0043]-¶[0049]; EX1004, ¶¶267-269.

E. Claim 14

Claim 14 is obvious over both Shieh and Swenson. EX1001, 11:61-63. Shieh discloses the use of "deep packet inspection (DPI)" as part of its analysis functionality. EX1006, ¶[0021]; *see id.*, ¶[0028], ¶[0040]-¶[0041]; EX1004, ¶¶270-273. Swenson also discloses that the flow analyzer in network controller 140 performs DPI. EX1007, ¶[0059]-¶[0060], ¶[0065].

F. Claim 15

Claim 15 is obvious in view of both Shieh and Swenson. EX1001, 11:64-67. As discussed above for Claim 14, Shieh discloses analysis of packets through DPI. EX1006, ¶[0021], ¶[0028], ¶[0040]-¶[0041]; EX1005, 3:11-12. A POSA would have understood that DPI refers to the inspection of at least part of the payload field. EX1004, ¶¶274-278. Further, Swenson discloses Claim 15 for the reasons discussed above in Ground 1.

G. Claim 16

Claim 16 is obvious over Lin in view of the knowledge of a POSA. EX1001, 12:1-4. A POSA would have known that a packet includes distinct header and payload fields. EX1008, Section 2; EX1004, ¶¶279-280. Further, Shieh states that the packet streams passing through its system include "TCP FIN or TCP RST packets." EX1006, ¶[0036]. A POSA would have understood that a TCP FIN packet includes a FIN flag bit in its header that is set. EX1004, ¶281. Further, a POSA would have understood that a TCP RST packet includes a RST flag bit in its header that is set. *Id.* A POSA would have found it obvious that the TCP FIN flag bit or the TCP RST flag bit being set could function as the packet-applicable criterion recited in Claim 1. *Id.*

H. Claim 17

Claim 17 is obvious over Shieh in view of the knowledge of a POSA. EX1001, 12:5-9. Shieh states that the packet streams passing through its system include TCP packets. EX1006, ¶[0027], ¶[0036]. Further, the TCP FIN or TCP RST packets discussed in Shieh render Claim 17 obvious for the reasons discussed above with respect to Claim 16. *Id.*, ¶[0036]; EX1004, ¶¶282-284. In addition, a POSA also would have understood that SYN and ACK were common flags in a TCP packet header that could have been used to identify packets in a packet stream. *Id.*

I. Claim 18

Claim 18 is obvious over Shieh in view of the knowledge of a POSA. EX1001, 12:10-16. Shieh discusses the header fields in a packet and how they are used to route packets. EX1005, [0044]. Similarly, Shieh discloses that IP or MAC addresses are packet-applicable criterion used to determine whether to forward packets to a security processing function based on whether the IP or MAC addresses match a predetermined address or addresses. EX1006, ¶[0031]; ¶[0027]; EX1004, ¶¶285-288.

J. Claim 19

Claim 19 is obvious over Shieh for the reasons as those discussed above for Claim 18. EX1001, 12:17-18; EX1006, ¶[0027]; ¶[0031]; EX1004, ¶289.

K. Claim 20

Claim 20 is obvious over Shieh in view of the knowledge of a POSA. EX1001, 12:20-27. Shieh discloses the transportation of TCP packets across a TCP network using TCP protocol. EX1006, ¶[0027], ¶[0036]; EX1004, ¶¶290-292. Further, Shieh describes tracking TCP packets and states that "a TCP/IP flow [i.e., a TCP packet] can be uniquely identified by … source and destination port…" EX1006, ¶[0027]. Shieh does not state that a TCP packet has a TCP sequence number or a TCP sequence mask field, but a POSA would have known those fields are always present in a TCP segment/packet. EX1004, ¶291. Further, Shieh teaches the use of a source

TCP port or a destination TCP port as the packet applicable criterion. EX1006, ¶[0027], ¶[0036]; EX1004, ¶291.

L. Claim 21

Claim 21 is obvious over Shieh. EX1001, 12:28-32. Shieh states that its packet network includes the Internet, a WAN or a LAN. EX1006, ¶[0020]. Shieh also discloses that its system can be in communication with an ISP. *Id.*, ¶[0053]; EX1004, ¶¶293-295.

M. Claim 22

Claim 22 is obvious over Shieh. EX1001, 12:33-36. Figure 1 of Shieh (reproduced below) shows that the source of the network packets (i.e., the claimed first entity) can be a server device (Nodes 206 or 207), and the destination of the claimed packets (i.e., the claimed second entity) can be a client device (Nodes 201 or 202), or vice versa.

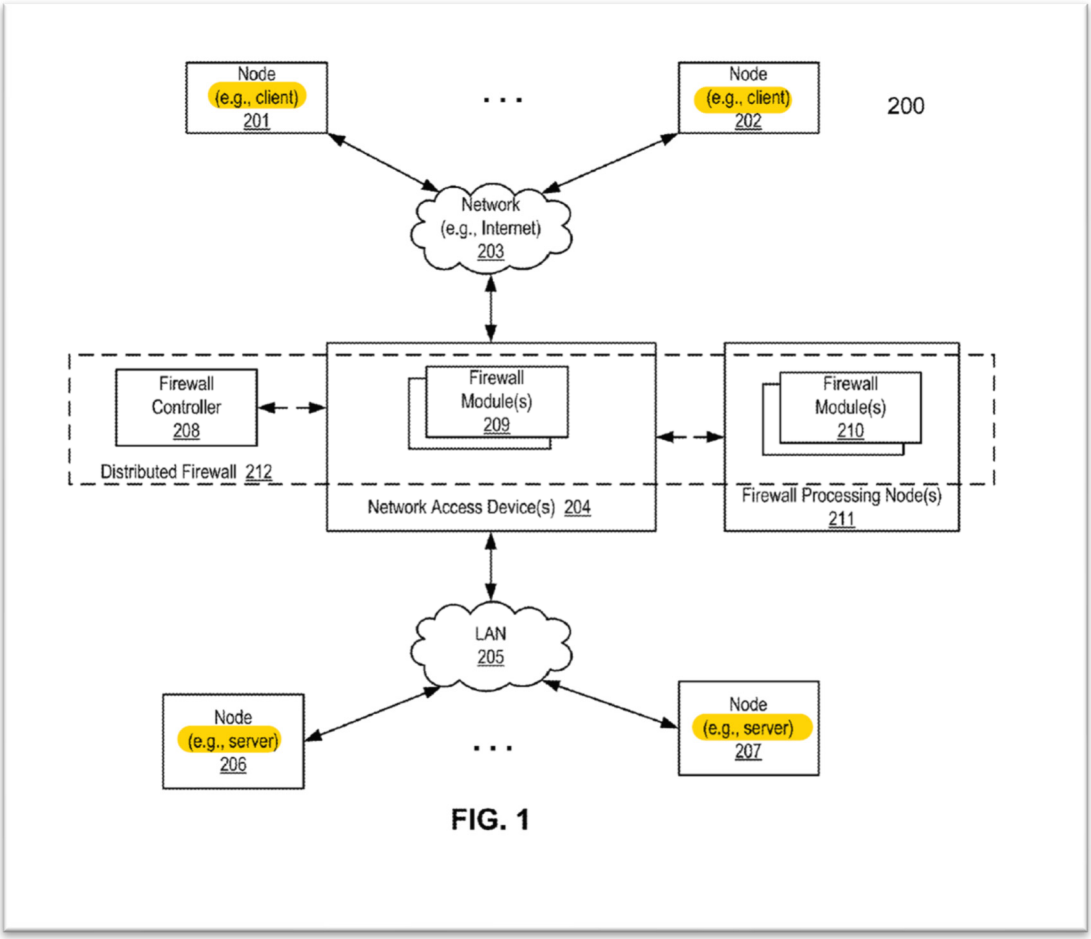


Figure 1 of Shieh (Annotated)

EX1006, Figure 1; *see id.*, ¶[0020], ¶[0022], ¶[0039], ¶[0053], ¶[0057], Figures 2A, 3; EX1004, ¶¶296-298.

N. Claim 23

Claim 23 is obvious over Shieh. EX1001, 12:37-41. As discussed above for Claim 22, Shieh discloses that the source or destination of the packets can be a server device, such as Nodes 206 or 207 in Figure 1. EX1006, Figure 1; *see id.*, ¶[0020], ¶[0022], ¶[0039], ¶[0053], ¶[0057], Figures 2A, 3. A POSA would have known that the server can be a web server. EX1004, ¶¶299-300; EX1006, ¶[0053]. Further,

Shieh discloses that "[a]ny of nodes 201-202 and 206-207 may be a client device (e.g., a desktop, laptop, Smartphone, gaming device) or a server." EX1006, ¶[0020]; EX1004, ¶¶301-302.

O. Claim 24

Claim 24 is obvious in view of Shieh. EX1001, 12:42-44. Shieh discloses the use of the standard OpenFlow protocol for communication between the network access devices (i.e., the network node) and the controller. EX1006, ¶[0025]-¶[0026]; EX1004, ¶303.

P. Claim 27

Claim 27 is obvious in view of Shieh. EX1001, 12:51-52. Shieh states that its network access devices (which correspond to the claimed network nodes) "may be a router or a gateway, a switch or an access point…" EX1006, ¶[0020]; *see id.*, ¶[0019], ¶[0025]-¶[0026], ¶[0032]; EX1004, ¶¶304-306.

Q. Claims 28

Claim 28 is obvious in view of Shieh. EX1001, 12:53-55. Shieh discloses the communication of IP packets through IP networks. EX1006, ¶[0038]; *see id.*, ¶[0027] (discussing the use of IP addresses to identify IP packets being transmitted across an IP packet network), ¶[0031] (same); EX1004, ¶¶307-309.

R. Claims 29

Claim 29 is obvious in view of Shieh. EX1001, 12:56-58. Shieh discloses the transportation of TCP packets across a TCP network using TCP protocol. EX1006, ¶[0027], ¶[0036]; EX1004, ¶¶310-311.

S. Claim 30

Claim 30 is obvious in view of Shieh. EX1001, 12:59-63. Shieh states that its method can be applied to "subsequent packets of a particular session." EX1006, ¶[0018]; *see id.*, ¶[0037]; EX1004, ¶¶312-314. Further, Shieh explains that its network access devices forward additional packets over the packet network in response to whether not the packets satisfy the packet-applicable criterion for the reasons discussed above for Elements [1.5]-[1.6]. EX1004, ¶¶315-318.

XI. OBJECTIVE INDICIA OF NONOBVIOUSNESS

At this stage of these proceedings, Petitioner has no burden to identify and rebut objective indicia of nonobviousness. EX1004, ¶319-320. Patent Owner must first present a *prima facie* case for such consideration, which Petitioner should then have the chance to rebut on reply. *Sega of Am., Inc. v. Uniloc USA, Inc.*, IPR2014-01453, Paper 11, at *20 (PTAB Mar. 10, 2015).

XII. DISCRETIONARY DENIAL UNDER § 325(D) OR § 314 IS NOT WARRANTED

Discretionary denial under 35 U.S.C. §325(d) is not warranted here because the challenges presented in this petition are neither cumulative nor redundant to the prosecution of the '111 Patent. The PTO has not previously considered Lin, Shieh or Swenson in connection with the claims of the '111 Patent. Further, the combination is not redundant to any combination of references considered during prosecution. Additionally, Dr. Bhattacharjee is an expert in the relevant art, and his analysis of the prior art has not been presented to the PTO. EX1004, ¶¶1-11. In view of the new information presented in this Petition, the Board has ample discretion to institute an IPR on the grounds presented.

Likewise, denial under 35 U.S.C. §314(a) is not warranted, as the application of the factors in *General Plastic Industrial Co., Ltd. v. Canon Kabushiki Kaisha* demonstrates that the filing of the instant Petition is not abusive, that none of the Patent Owner, the PTO, or the Board has addressed the merits of the grounds in this Petition, and instituting the present proceeding is an efficient use of the Board's resources. *See* IPR2016-01357, Paper 19 (PTAB Sept. 6, 2017) (precedential). Petitioner has not challenged the '111 Patent in a prior petition.

Moreover, the Petition should not be denied under the discretionary factors set out in *Apple Inc. v. Fintiv, Inc.,* IPR2020-00019, Paper 11, at 6 (PTAB March 20, 2020). Petitioner addresses each of those factors below:

1. Whether the court granted a stay – A motion to stay has not yet been filed in the parallel litigation, so the Board should not infer 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 (PTAB June 16, 2020) (informative); *Dish Network LLC v. Broadband iTV, Inc.*, IPR2020-01359, Paper 15 at 11 (PTAB Feb. 12, 2021). Thus, this factor is neutral on discretionary denial.

2. Parallel proceeding trial date – This factor weighs against discretionary denial because the projected trial date in the parallel litigation is "around the same time" as the Board's expected Final Written Decision. The parallel litigation was filed in the U.S. District Court for the Eastern District of Texas on July 22, 2022. EX1010. While trial is currently proposed for March 4, 2024 (EX1011 at 1), the Board recognizes "that scheduled trial dates are unreliable and often change." EX1012 at 8. For this reason, the Board now uses median time-to-trial statistics in the litigation venue to determine a projected trial date. EX1012 at 8-9. The median time-to-trial in the Eastern District of Texas was 24.5 months as of June 30, 2022. EX1013. Thus, the projected trial date in the litigation for *Fintiv* purposes is August 2024, approximately 24.5 months after July 2022. The Board's Final Written

Decision is also expected in August 2024, which is 18 months after Petitioner expects a notice of accorded filing for this Petition. Accordingly, this factor weighs against discretionary denial.

3. Investment in the parallel proceeding by the court and the parties – The parallel litigation is in its early stages, and the Court has not issued any substantive ruling relating to the '111 Patent. The Parties have not exchanged preliminary positions on claim construction, and the claim construction hearing is not scheduled until September 7, 2023, a month after the expected institution decision by the Board. EX1011 at 3-4. Further, the Parties only exchanged preliminary invalidity contentions on February 2, 2023, 19 days before this Petition was filed. *Id.* at 5. Moreover, the Parties have not exchanged their first set of fact discovery requests, and expert discovery does not begin until October 19, 2023. *Id.* at 3. The early stage of and minimal investment in the parallel litigation weighs against discretionary denial. *See PEAG LLC v. Varta Microbattery GMBH*, IPR2020-01214, Paper 8 at 17 (Jan. 6, 2021).

4. Overlapping issues with the parallel litigation proceeding – Preliminary invalidity contentions were only served in the parallel litigation on February 2, 2023, and thus it is too early to determine overlapping invalidity issues. EX1011 at 5. Nonetheless, instituting a proceeding will allow the Board to address the art, and a Final Written Decision would narrow the issues in the parallel litigation due to the

estoppel provisions of 35 U.S.C. § 315(e)(2). Moreover, there will be no overlap of prior art issues because, if the Board institutes trial, Petitioner will cease asserting in the parallel litigation the combination of references on which trial is instituted for the claims on which trial is instituted, to the extent Petitioner even asserts the same combination in the parallel litigation. 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. Identity of the Parties – Petitioner is a defendant in the parallel litigation, but that is true of most petitioners in IPR proceedings. Accordingly, this factor is neutral on discretionary denial.

6. Other circumstances, including the merits – As discussed in detail above, the analysis in Grounds 1-2 provides a compelling unpatentability challenge to Claims 1-9, 12-24 and 27-31. The merits of Petitioners' arguments are strong, and this factor weighs against discretionary denial. *See* EX1012 at 3-5; *Sand Revolution II*, IPR2019-01393, Paper 24 at 13.

XIII. MANDATORY NOTICES UNDER 37 C.F.R. § 42.8

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

The '111 Patent is asserted in *Orckit Corp. v. Cisco Systems, Inc.*, Case No. 2:22-cv-00276 (E.D. Tex.).

C. Lead and Back-up Counsel and Service Information

Lead Counsel	
Jeffrey D. Blake MERCHANT & GOULD P.C. 191 Peachtree Street NE Suite 3800 Atlanta, GA 30303	Phone: 404-954-5040 Fax: 612-332-9081 jblake@merchantgould.com USPTO Reg. No. 58,884
Back-up Counsel	
Daniel W. McDonald MERCHANT & GOULD P.C. 150 South Fifth Street Suite 200 Minneapolis, MN 55402	Phone: 612-336-4637 Fax: 612-332-9081 dmcdonald@merchantgould.com USPTO Reg. No. 32,044

Please address all correspondence to lead and back-up counsel at OrckitIPR@merchantgould.com. Petitioner consents to electronic service.

XIV. CONCLUSION

For the reasons above, Petitioners ask that the Board order an *inter partes* review trial for Claims 1-9, 12-24 and 27-31, and that the Director cancel these claims as unpatentable.

Respectfully submitted,

Date: February 21, 2023

/Jeffrey D. Blake/
Jeffrey D. Blake
Registration No. 58,884

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24, the undersigned attorney for the Petitioner, Cisco Systems, Inc., declares that the argument section of this Petition has 13,982 words, according to the word count tool in Microsoft Word™.

/Jeffrey D. Blake/
Jeffrey D. Blake
Registration No. 58,884

APPENDIX A – CLAIM LISTING

U.S. Patent No. 10,652,111

Claim or Element #	Claim Language
Claim 1	
[1.0]	A method for use with a packet network including a network node for transporting packets between first and second entities under control of a controller that is external to the network node, the method comprising:
[1.1]	sending, by the controller to the network node over the packet network, an instruction and a packet-applicable criterion;
[1.2]	receiving, by the network node from the controller, the instruction and the criterion;
[1.3]	receiving, by the network node from the first entity over the packet network, a packet addressed to the second entity;
[1.4]	checking, by the network node, if the packet satisfies the criterion;
[1.5]	responsive to the packet not satisfying the criterion, sending, by the network node over the packet network, the packet to the second entity; and
[1.6]	responsive to the packet satisfying the criterion, sending the packet, by the network node over the packet network, to an entity that is included in the instruction and is other than the second entity.
Claim 2	
	The method according to claim 1, wherein the instruction is 'probe', 'mirror', or 'terminate' instruction, and upon receiving by the network node the 'terminate' instruction, the method further comprising blocking, by the network node, the packet from being sent to the second entity and to the controller.
Claim 3	
	The method according to claim 1, wherein the instruction is a 'probe', a 'mirror', or a 'terminate' instruction, and upon receiving by the network node the 'mirror' instruction and responsive to the packet satisfying the criterion, the method

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 314 of 496

Claim or Element #	Claim Language
	further comprising sending the packet, by the network node, to the second entity and to the controller.
Claim 4	
	The method according to claim 1, wherein the instruction is 'probe', 'mirror', or 'terminate' instruction, and upon receiving by the network node the 'probe' instruction and responsive to the packet satisfying the criterion, the method further comprising: sending the packet, by the network node, to the controller; responsive to receiving the packet, analyzing the packet, by the controller; sending the packet, by the controller, to the network node; and responsive to receiving the packet, sending the packet, by the network node, to the second entity.
Claim 5	
	The method according to claim 1, further comprising responsive to the packet satisfying the criterion and to the instruction, sending the packet or a portion thereof, by the network node, to the controller.
Claim 6	
	The method according to claim 5, further comprising storing the received packet or a portion thereof, by the controller, in a memory.
Claim 7	
	The method according to claim 5, further comprising responsive to the packet satisfying the criterion and to the instruction, sending a portion of the packet, by the network node, to the controller.
Claim 8	
	The method according to claim 7, wherein the portion of the packet consists of multiple consecutive bytes, and wherein the instruction comprises identification of the consecutive bytes in the packet.
Claim 9	
	The method according to claim 5, further comprising responsive to receiving the packet, analyzing the packet, by the controller.

Claim or Element #	Claim Language
Claim 12	
	The method according to claim 9, wherein the analyzing comprises applying security or data analytic application.
Claim 13	
	The method according to claim 9, wherein the analyzing comprises applying security application that comprises firewall or intrusion detection functionality.
Claim 14	
	The method according to claim 9, wherein the analyzing comprises performing Deep Packet Inspection (DPI) or using a DPI engine on the packet.
Claim 15	
	The method according to claim 9, wherein the packet comprises distinct header and payload fields, and wherein the analyzing comprises checking part of, or whole of, the payload field.
Claim 16	
	The method according to claim 1, wherein the packet comprises distinct header and payload fields, the header comprises one or more flag bits, and wherein the packet-applicable criterion is that one or more of the flag bits is set.
Claim 17	
	The method according to claim 16, wherein the packet is an Transmission Control Protocol (TCP) packet, and wherein the one or more flag bits comprises comprise a SYN flag bit, an ACK flag bit, a FIN flag bit, a RST flag bit, or any combination thereof.
Claim 18	
	The method according to claim 1, wherein the packet comprises distinct header and payload fields, the header comprises at least the first and second entities addresses in the packet network, and wherein the packet-applicable criterion is that the first entity address, the second entity address, or both match a predetermined address or addresses.

Claim or Element #	Claim Language
Claim 19	
	The method according to claim 18, wherein the addresses are Internet Protocol (IP) addresses.
Claim 20	
	The method according to claim 1, wherein the packet is an Transmission Control Protocol (TCP) packet that comprises source and destination TCP ports, a TCP sequence number, and a TCP sequence mask fields, and wherein the packet-applicable criterion is that the source TCP port, the destination TCP port, the TCP sequence number, the TCP sequence mask, or any combination thereof, matches a predetermined value or values.
Claim 21	
	The method according to claim 1, wherein the packet network comprises a Wide Area Network (WAN), Local Area Network (LAN), the Internet, Metropolitan Area Network (MAN), Internet Service Provider (ISP) backbone, datacenter network, or inter-datacenter network.
Claim 22	
	The method according to claim 1, wherein the first entity is a server device and the second entity is a client device, or wherein the first entity is a client device and the second entity is a server device.
Claim 23	
	The method according to claim 22, wherein the server device comprises a web server, and wherein the client device comprises a smartphone, a tablet computer, a personal computer, a laptop computer, or a wearable computing device.
Claim 24	
	The method according to claim 22, wherein the communication between the network node and the controller is based on, or uses, a standard protocol.
Claim 27	
	The method according to claim 1, wherein the network node comprises a router, a switch, or a bridge.

Claim or Element #	Claim Language
Claim 28	
	The method according to claim 1, wherein the packet network is an Internet Protocol (IP) network, and the packet is an IP packet.
Claim 29	
	The method according to claim 28, wherein the packet network is an Transmission Control Protocol (TCP) network, and the packet is an TCP packet.
Claim 30	
[30.0]	The method according to claim 1, further comprising: receiving, by the network node from the first entity over the packet network, one or more additional packets; checking, by the network node, if any one of the one or more additional packets satisfies the criterion;
[30.1]	responsive to an additional packet not satisfying the criterion, sending, by the network node over the packet network, the additional packet to the second entity; and responsive to the additional packet satisfying the criterion, sending the additional packet, by the network node over the packet network, in response to the instruction.
Claim 31	
	The method according to claim 1, wherein the packet network is a Software Defined Network (SDN), the packet is routed as part of a data plane and the network node communication with the controller serves as a control plane.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Cisco Systems, Inc.	§	Petition for *Inter Partes* Review
	§	
Petitioner	§	U.S. Patent No. 10,652,111
	§	
	§	

CERTIFICATE OF SERVICE

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

Date of service	February 21, 2023
Manner of service	PRIORITY MAIL EXPRESS
Documents served	Petition for *Inter Partes* Review, including Exhibit List; Exhibits 1001 through 1013; Power of Attorney
Persons served	May Patents Ltd. c/o Dorit Shem-Tov P.O.B. 7230 Ramat-Gan, 5217102 Israel

/Jeffrey D. Blake/
Jeffrey D. Blake
Registration No. 58,884

EXHIBIT 4

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner

v.

ORCKIT IP, LLC.
Patent Owner

IPR2023-00714
U.S. Patent No. 6,680,904

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

TABLE OF CONTENTS

I. GROUNDS FOR STANDING 6

II. SUMMARY OF THE '904 PATENT 6

A. Overview 6

B. Prosecution History 11

III. LEVEL OF ORDINARY SKILL IN THE ART 12

IV. CLAIM CONSTRUCTION 13

A. "second master unit" 13

V. RELIEF REQUESTED 15

VI. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE 15

A. Challenged Claims And Statutory Grounds For Challenge 15

B. Grounds 1-2: Claims 1-26 Are Obvious Over Vink With Patrick (Ground 1) Or Over Vink With Patrick And AAPA (Ground 2). 19

1. Independent Claims 1, 4, 9, and 11 19

2. Dependent Claim 5 65

3. Dependent Claim 6 70

4. Dependent Claim 7 71

5. Dependent Claim 8 76

6. Dependent Claim 10 76

7. Dependent Claim 12 78

8. Dependent Claim 13 83

10. Dependent Claim 15 87

11. Dependent Claim 16 87

12. Dependent Claim 17 89

13. Dependent Claim 18 89

14. Independent Claim 19 91

15. Dependent Claim 20 93

16. Dependent Claim 21 94

17. Dependent Claim 22 94

18. Dependent Claim 23 95

19. Dependent Claim 24 95

20. Dependent Claims 2 and 25 96

21. Dependent Claims 3 and 26 98

VII. DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE.....99

A. Discretionary Denial Under *Fintiv* Is Not Appropriate 99

B. Discretionary Denial Under 35 U.S.C. § 325(d) Is Not Appropriate..... 100

C. Discretionary Denial Under *General Plastic* Is Not Appropriate.... 101

VIII. CONCLUSION.....101

IX. MANDATORY NOTICES 102

A. Real Party-in-Interest 102

B. Related Matters..... 102

C. Lead and Back-up Counsel and Service Information 102

PETITIONER'S EXHIBIT LIST

1001	U.S. Patent No. 6,680,904 to Cohn et al. ("'904 patent)
1002	Prosecution History of U.S. Patent No. 6,680,904
1003	Declaration of Dr. Henry Houh
1004	*Curriculum Vitae* of Dr. Houh
1005	International Patent Publication No. WO 91/14324 ("Vink")
1006	U.S. Patent No. 5,790,541 ("Patrick")
1007	U.S. Patent No. 6,631,136 ("Chowdhury")
1008	U.S. Patent No. 5,425,026 ("Mori")
1009	ATM Volume I: Foundation for Broadband Networks
1010	First Amended Docket Control Order, *Orckit Corporation v. Cisco Systems, Inc*., Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, January 15, 2023)
1011	Scheduling Order in *Orckit Corp. v. Cisco Systems, Inc.*, Case No. 2:22-cv-00276 (E.D. Tex. October 18, 2022)

*** Unless indicated otherwise, all annotations in figures and emphases in quoted material have been added throughout the Petition.

Pursuant to 35 U.S.C. §§ 311, 314(a), and 37 C.F.R. § 42.100, Cisco Systems, Inc. ("Petitioner") respectfully requests cancellation of claims 1-26 ("Challenged Claims") of U.S. Patent No. 6,680,904 ("'904 patent," Ex-1001) as unpatentable under (pre-AIA) 35 U.S.C. § 103(a).

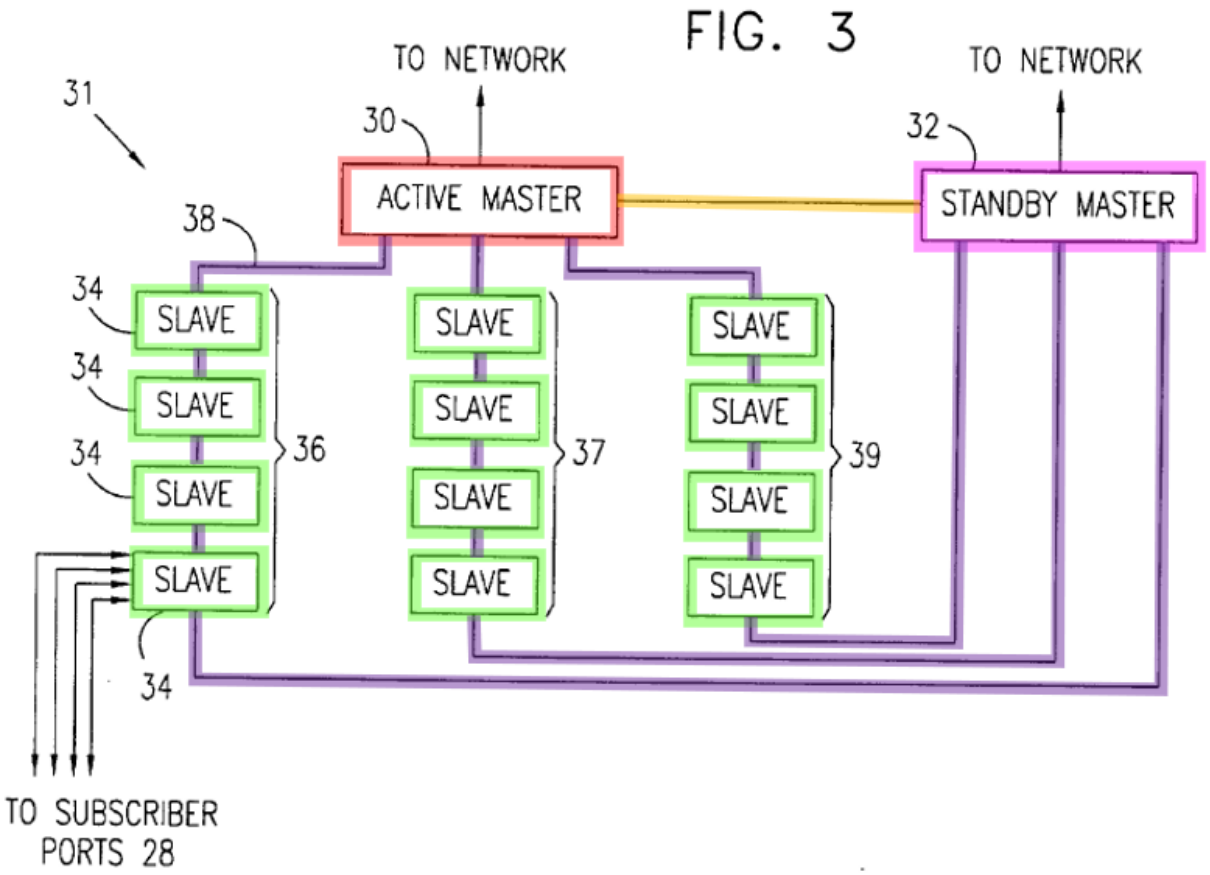
I. GROUNDS FOR STANDING

Petitioner certifies that the '904 patent is eligible for IPR and that Petitioner is not barred or estopped from requesting IPR challenge. 37 C.F.R. § 42.104(a).

II. SUMMARY OF THE '904 PATENT

A. Overview

The '904 patent relates to "high speed data communication systems." Ex-1001, Abstract. Figure 3, below, illustrates the disclosed network access system topology:

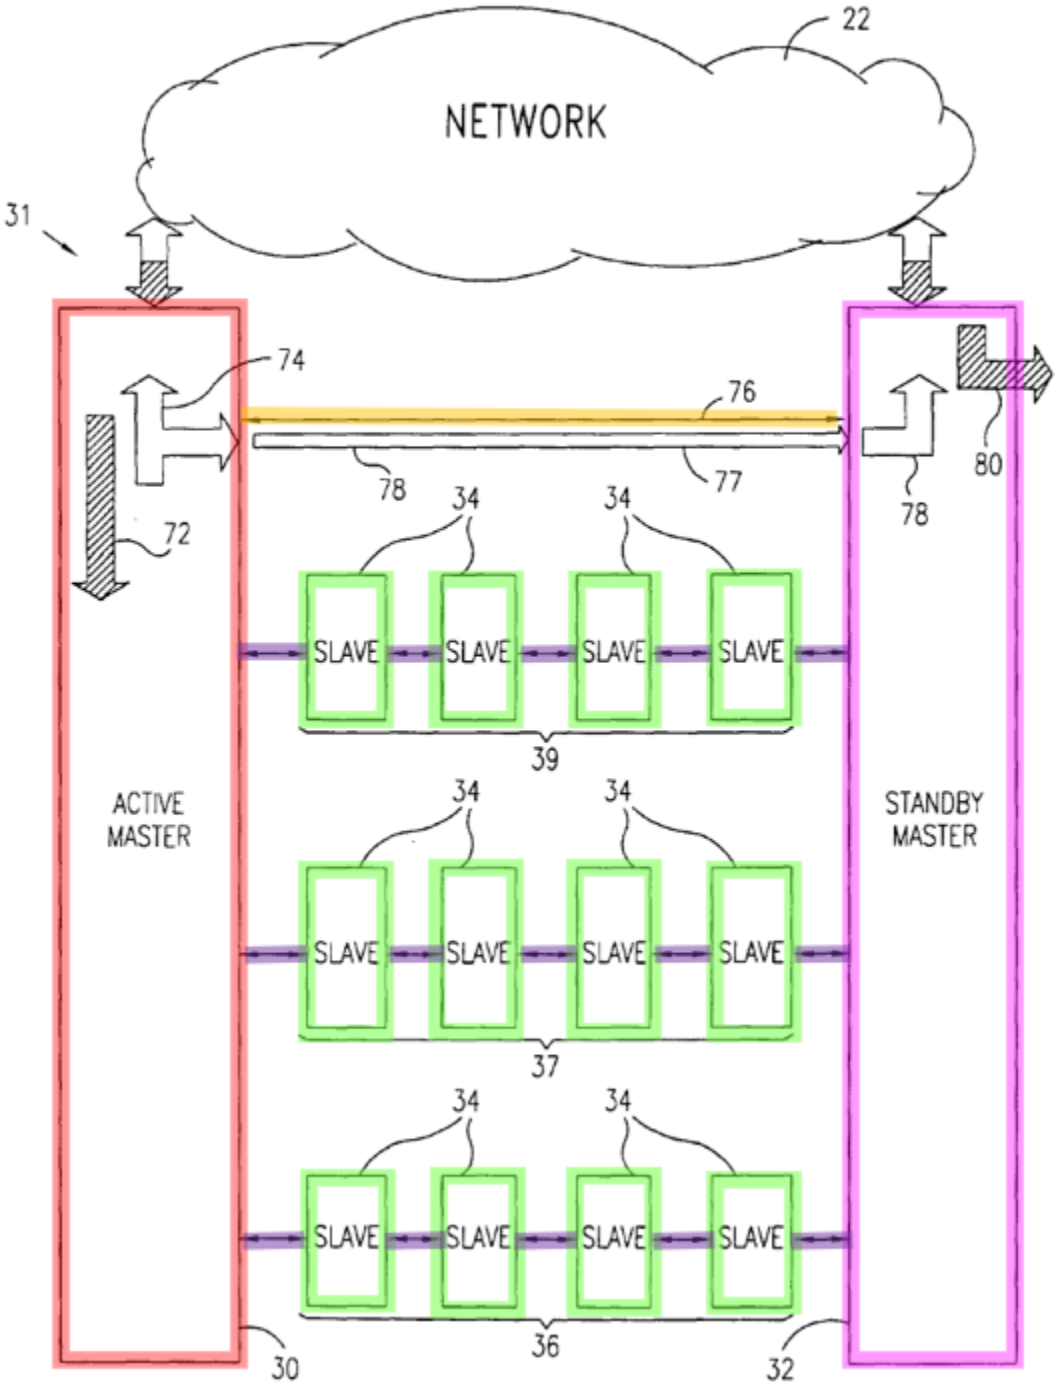


Ex-1001, Figure 3. As shown, a "plurality of **slaves 34** are connected between **master[] 30** and **[master] 32** by **lines 38** in a number of daisy chains 36, 37, and 39." Ex-1001, 6:26-28. "Each **slave** comprises multiple subscriber ports 28, which link system 31 to respective subscriber locations." *Id.*, 6:28-29.

Under normal conditions (Figure 5 below), downstream data flows down from network 22 to **slaves 34** in each of daisy chains 36, 37, and 39 by passing through **active master 30** (hatched arrow 72). In contrast, the downstream data passing through **standby master 32** is discarded (hatched arrow 80). *See* Ex-1001,

8:31-41. Moreover, upstream data (data flowing from the **slaves** to network 22) is passed through the **active master 30**, which sends it up to network 22 and also to **standby master 32** (open arrow 77) over **protection interface 76**, which in turn also sends the upstream data up to the network (open arrow 78). *See id.*, 8:42-49. "This redundancy in transmission is in accordance with fault protection mechanism … known in the art." *Id.*, 8:49-52.

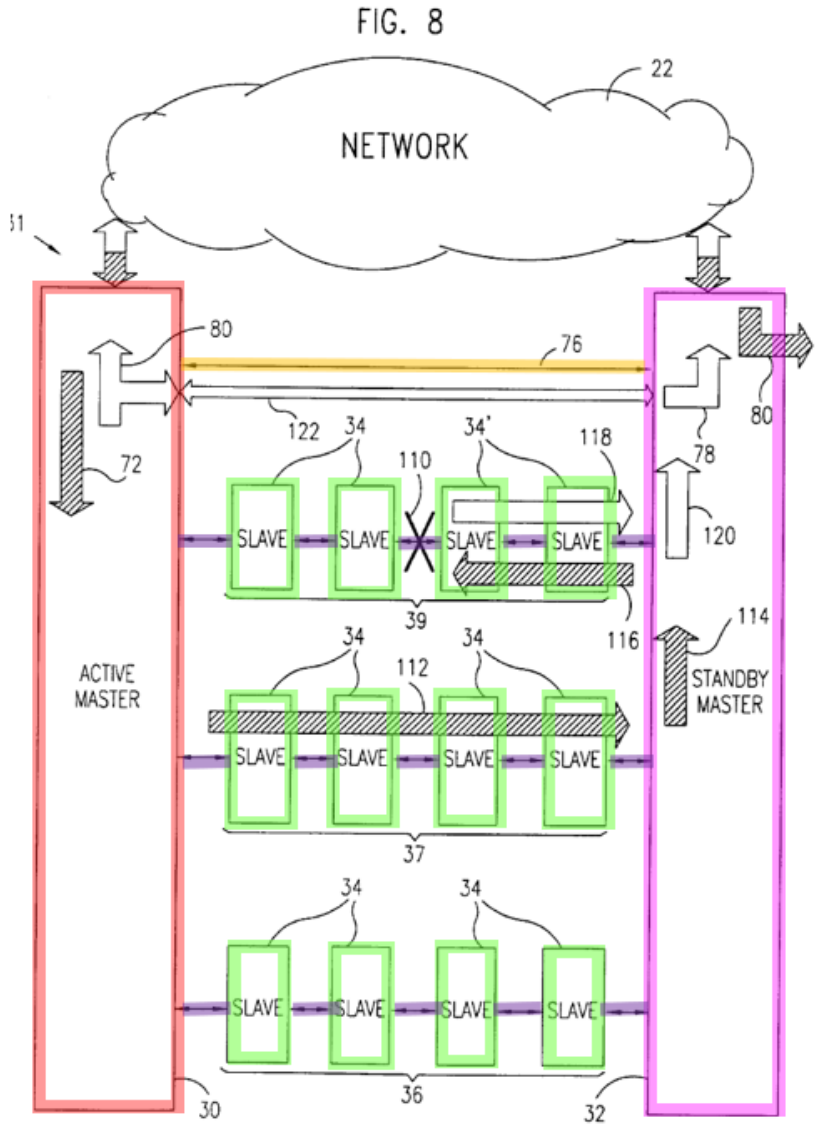
FIG. 5



When a fault occurs in one daisy chain 39 (see "X" in Figure 8 below),

active master 30 reroutes downstream data packets to the next daisy chain 37

(dashed arrow 112), and **standby master 32** switches that traffic back onto daisy chain 39 (dashed arrow 116) towards **slaves 34'**. *See* Ex-1001, 9:38-57. Upstream data from **slaves 34'** is similarly routed to **standby master 32** (open arrow 118). *See id.*, 9:57-65.

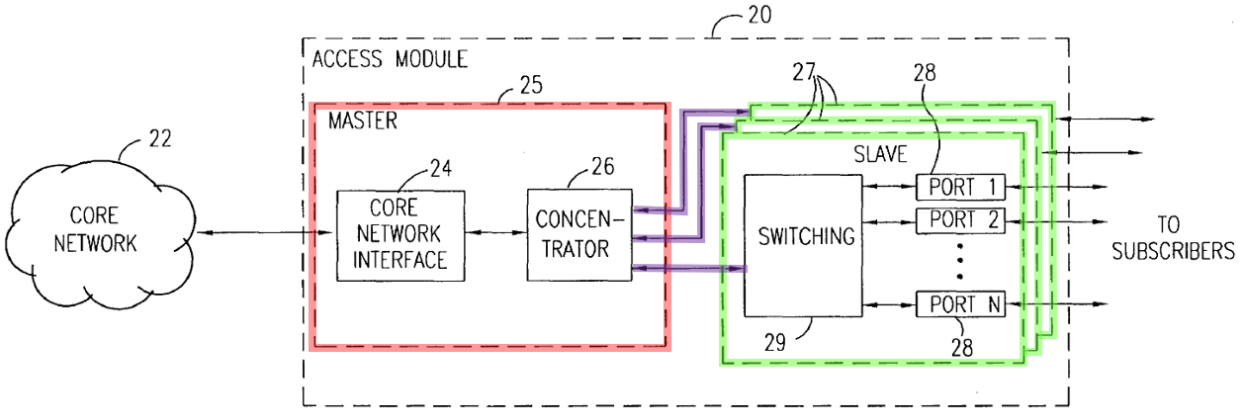


Ex-1001, Figure 8; Ex-1003, ¶¶34-37.

B. Prosecution History

The '904 Patent was filed on December 27, 1999, with 26 original claims. Ex-1002, 28-91. The Examiner rejected claims 1 and 12 over Applicant Admitted Prior Art ("AAPA"), citing Figures 1 and 2B. *Id.*, 158-160.

FIG. 1
PRIOR ART



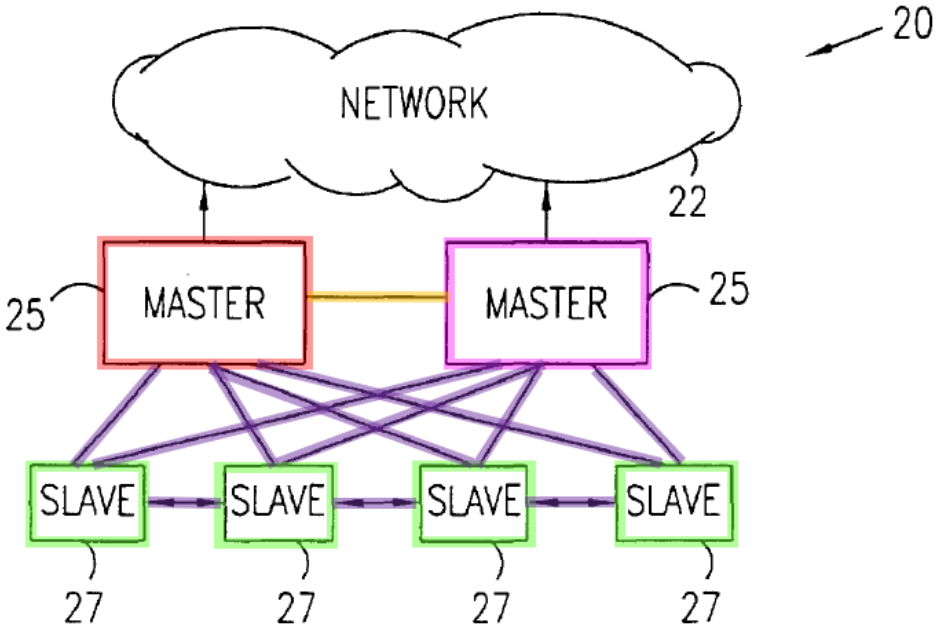


FIG. 2B
PRIOR ART

The Examiner further rejected dependent claim 13 by combining AAPA with U.S. Patent No. 6,181,715 ("Phillips"), holding that it would have been obvious to employ a DSLAM unit connecting to an IP network. *Id.*, 160-161. The Examiner held that claims 14 and 19-26 are allowed and that claims 2-11 and 15-18 would be allowable. *Id.,* 161.

In response, Applicant amended claim 1 to recite "a second slave unit connected to the first slave unit but not connected to the first or second master unit" (Ex-1002, 172) and amended claims 2, 7 and 9, which issued as independent claims 4, 9 and 11. *Id.,* 172-175. Original independent claims 14 and 19 were not amended. *Id.*, 176-179. The patent issued on January 20, 2004. Ex-1001, 1; Ex-1003, ¶¶38-46.

III. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill In The Art ("POSITA") would have had a bachelor's degree in computer science, computer engineering, electrical engineering, or equivalent training, or approximately two years of experience working in the field of information technology and networking as of December 27, 1999. Lack of professional experience can be substituted by additional education, and vice versa. Ex-1003, ¶¶59-60.

IV. CLAIM CONSTRUCTION

Claim terms in IPRs are construed according to their "ordinary and customary meaning" to those of skill in the art. 37 C.F.R. § 42.100(b).

A. "second master unit"

Petitioner submits that, for the purposes of this proceeding, the claim term "second master unit" should be construed as "a standby master that receives upstream packets from the first master unit during normal operations."

The term "second master unit" appears in claims 1-13 and 19-26. These claims additionally recite a "first master unit" and "slave units," which presumably have different meaning. *See Bancorp Servs., LLC v. Hartford Life Ins. Co.,* 359 F.3d 1367, 1373 (Fed. Cir. 2004) ("[U]se of both terms in close proximity in the same claim gives rise to an inference that a different meaning should be assigned to each."). However, because the Board "need only construe terms … only to the extent necessary to resolve the controversy," construing the term "second master unit" is sufficient to provide clarity to the different meanings assigned to these three terms. *See Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017).

Beginning with the claims, the claims do not explicitly assign different meanings to "second master unit" and "first master unit." For example, the independent claims 1, 4, 9 and 11 merely require the "first master unit" and the

"second master unit" to each have "a physical interface to a packet-switched network."

Turning to the specification, the '904 patent describes "two mutually-linked masters 30 and 32." Ex-1001, 6:22-26. "Master 30 is termed the active master," and master 32 is termed the "standby master 32." Ex-1001, 6:38-44. While the dependent claims suggest that the "second master unit" corresponds to the standby master, (*see*, *e.g.*, claims 5-8, 12, and 20-23), the specification notes that (1) the active and standby masters are "structurally substantially identical" to each other, (Ex-1001, 9:8-9), and (2) the standby master can be "treated as an additional slave on one of the chains." Ex-1001, 9:24-27.

Hence, properly assigning different meanings to the claim terms requires ascertaining the function performed by the standby master. The standby master's function that distinguishes it from the active master ("first master unit") and the "slave units" is the ***receipt of upstream packets from the active master ("first master unit") during normal operations***. *See* Ex-1001, 8:44-47 ("The active master… send[s] [these packets]… to standby master 32…."); *see also id.*, 2:23-24 ("upstream packets are bicast by the active master to both the core network and to the standby master."), 6:43-44 ("Active master 30 bicasts upstream packets to the network and to standby master 32."), 9:16-22 ("Upstream data packets received by the standby master… from active master 30…."). To be clear, other functions

are performed by the standby master during normal operations, such as "transmit[ting] the upstream packets to the core network." Ex-1001, 2:25-26. But that function is performed by both the active and standby master. The function of receiving, from the active master, upstream packets that is sent by the slave units during normal operation is only performed by the standby master. Hence, to provide clarity to the claim terms, the "second master unit" should be construed as "a standby master that receives upstream packets from the first master unit during normal operations." Ex-1003, ¶¶61-71.

V. RELIEF REQUESTED

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

VI. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

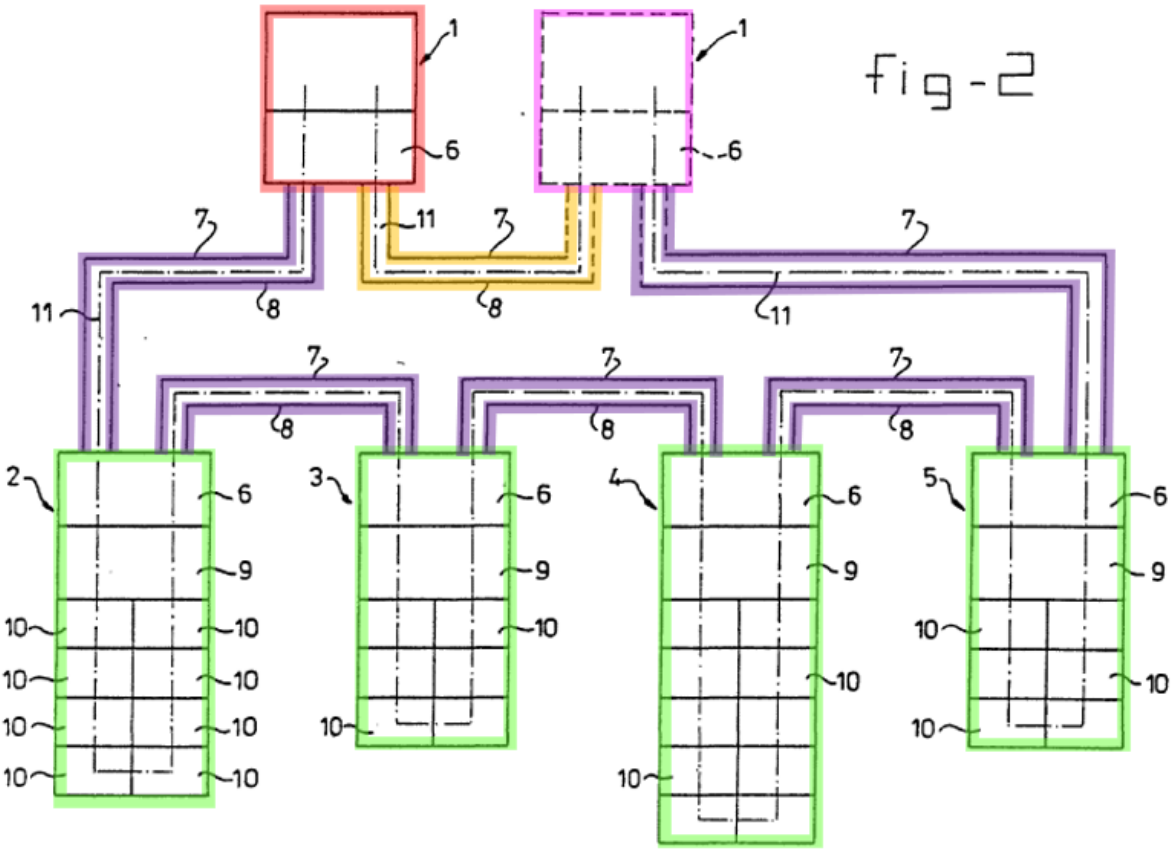
A. Challenged Claims And Statutory Grounds For Challenge

Grounds	Claims	Basis
#1	1-26	Vink with Patrick
#2	1-26	Vink with Patrick and AAPA

1. Vink

WO 91/14324 ("Vink," Ex-1005) published on September 19, 1991, and qualifies as prior art under §102(b).

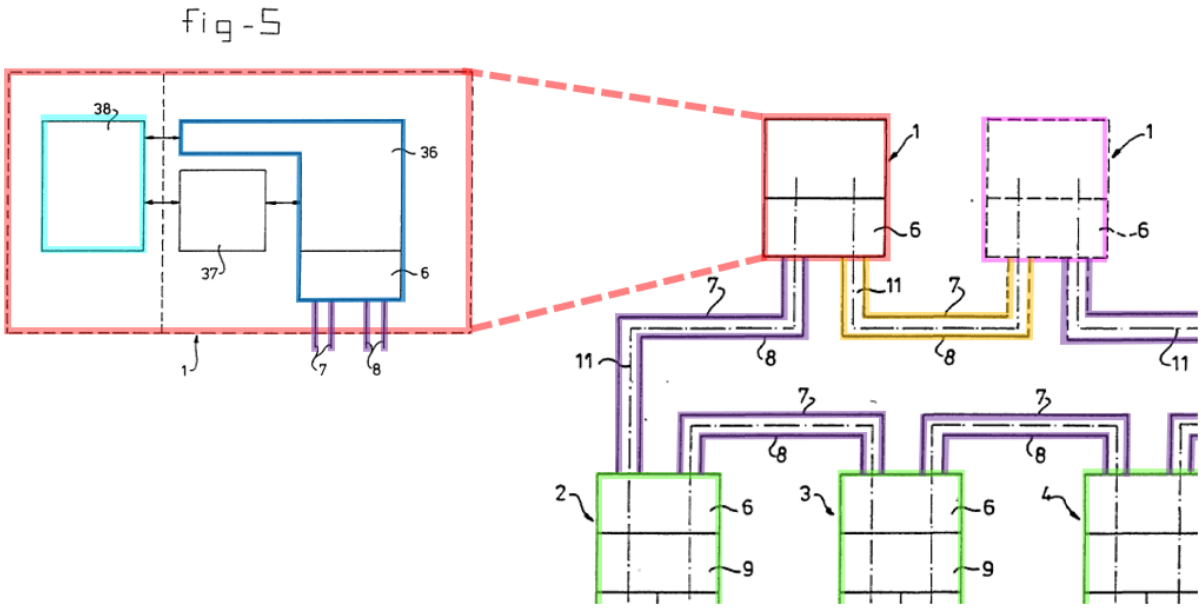
As shown below (Figure 2), Vink is directed to a communication system that exchanges data across a network in serial form under the control of at least one **master station**. Ex-1005, 1. Vink's communication system has two **master stations** and **substations (slaves 2, 3, 4, and 5)** that are connected to each other in series via the **transmission lines 7, 8**. *See* Ex-1005, 15:28-35. The two **master stations** are also connected to each other via **transmission lines 7, 8**.



Ex-1005, Figure 2.

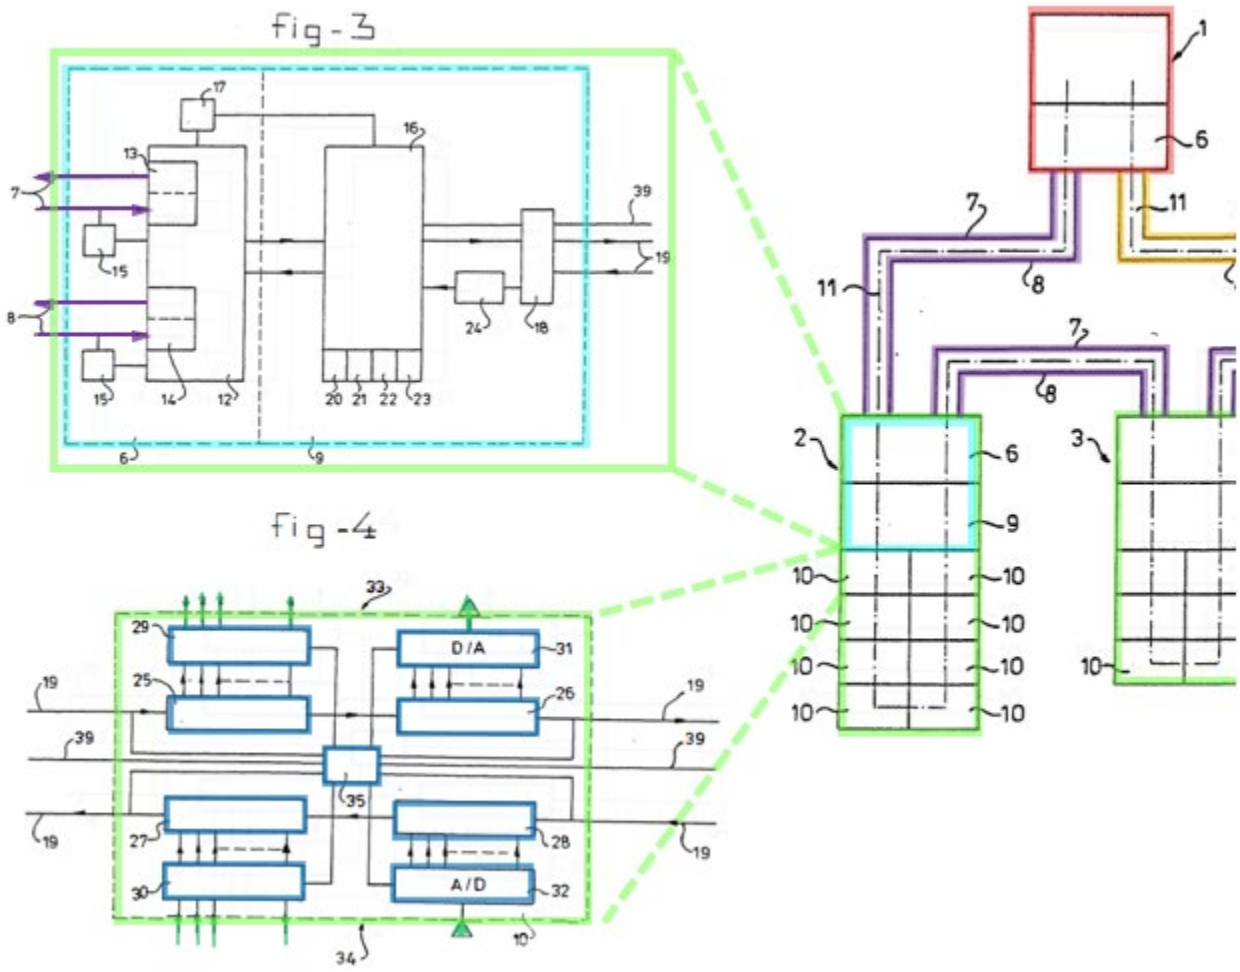
Figures 3-4 and 5 provide additional details regarding the **slave** and **master** components, respectively. As shown below, the **master** includes

transmitting/receiving unit 6, control unit 36, memory 37, and **control computer 38**.



Ex-1005, Figures 2 and 5. **Control computer 38** determines the slaves to/from which data packets should be delivered/are received from. *See* Ex-1005, 21:4-7. **Control unit 36 and transmission/receiving unit 6** transmit/receive the data packets via **transmission lines 7, 8** as dictated by **control computer 38**. *See* Ex-1005, 20:30-21:4.

As shown below, each **slave** contains **transmitting/receiving unit 6**, **control module 9** (Figure 3), and one or more **I/O modules** (Figure 4).



Ex-1005, Figures 2-4. The **transmitting/receiving unit 6** transmits/receives data packets to/from **transmission lines 7, 8** and the components within **transmitting/receiving unit 6** and **control module 9** determine whether the current slave should receive the data packets. *See* Ex-1005, 11:38-12:4, 14:12-14, 25:8-24. If the data packets are intended for a next subsequent slave, they are forwarded onto the next slave. *See* Ex-1005, 17:3-18:33. If the data packets are intended for one of the **I/O modules** in the current slave, they are passed serially

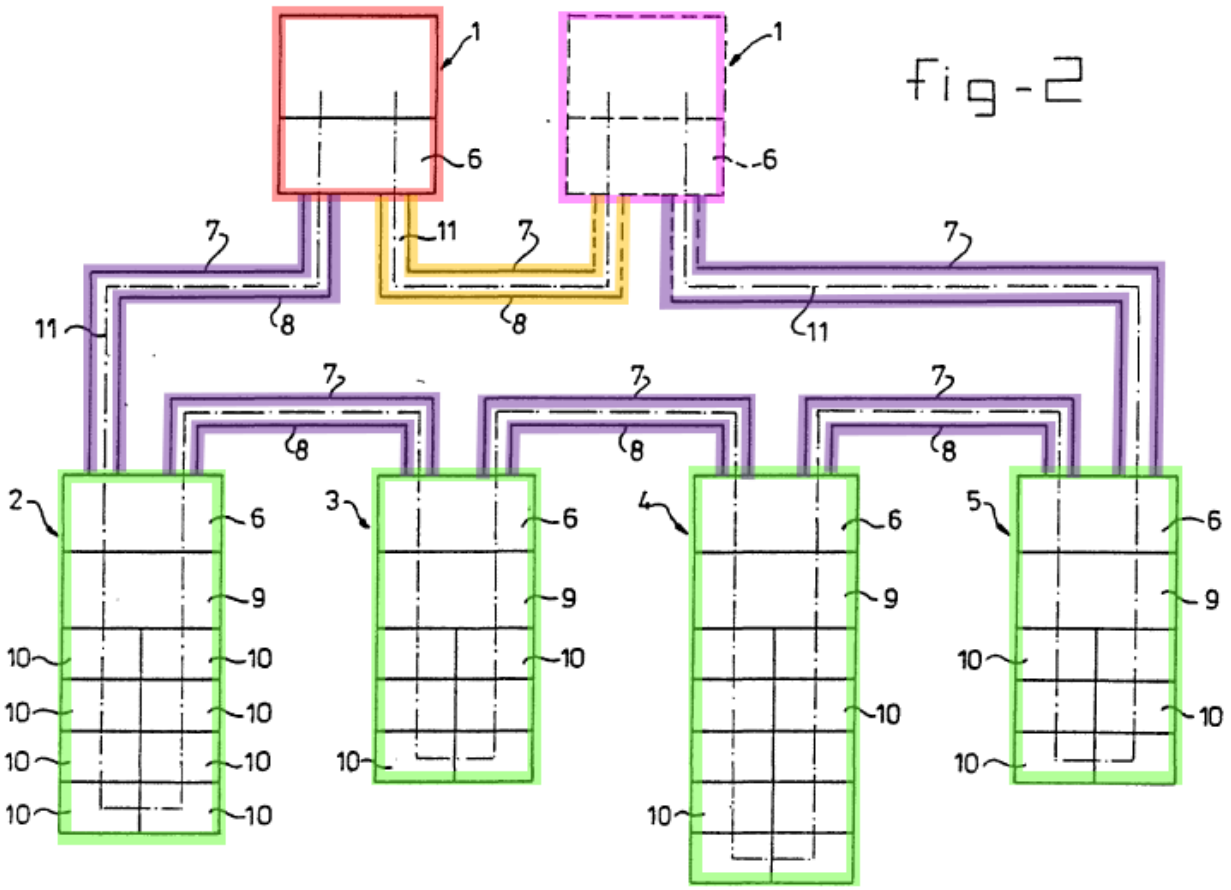
onto the **I/O modules** and **control means 35**, the **shift registers 25, 26, 30, and 32**, and **buffers 29 and 30** in the correct I/O modules route the data packets to **output side 33**. *See* Ex-1005, 19:1-20:29, 7:9-16, 23:14-21. These **I/O modules** further include **input side 34** that accepts data packets from the devices connected to the **I/O module**. *See id.*

B. Grounds 1-2: Claims 1-26 Are Obvious Over Vink With Patrick (Ground 1) Or Over Vink With Patrick And AAPA (Ground 2).

1. Independent Claims 1, 4, 9, and 11

a) [1pre/4pre/9pre/11pre] Network access apparatus, comprising:

To the extent limiting, Vink with Patrick (and AAPA for Ground 2) suggests the preambles. Vink discloses "a communication system [e.g., Local Area Networks (LANs)] for exchanging data … under the control of at least one **master station** (master)." Ex-1005, 1:6-15. As shown below, a plurality of **substations (slaves) 2, 3, 4, 5** are connected in series to two different **master units**. *See* Ex-1005, Abstract. Ex-1003, ¶¶74-75.



Ex-1005, Figure 2.

Although Vink does not explicitly teach that **master 1** accesses a network, **master stations** connected to networks were well-known. For example, as shown below, Patrick discloses a "communication system … characterized by a topology having a **primary node** connected to a first network, such as the Internet." Ex-1006, Abstract. And although only a single **primary station** is depicted in Figure 1, Patrick teaches "**more than one primary station**." Ex-1006, 3:27-32. *See* Ex-1003, ¶¶76-77.

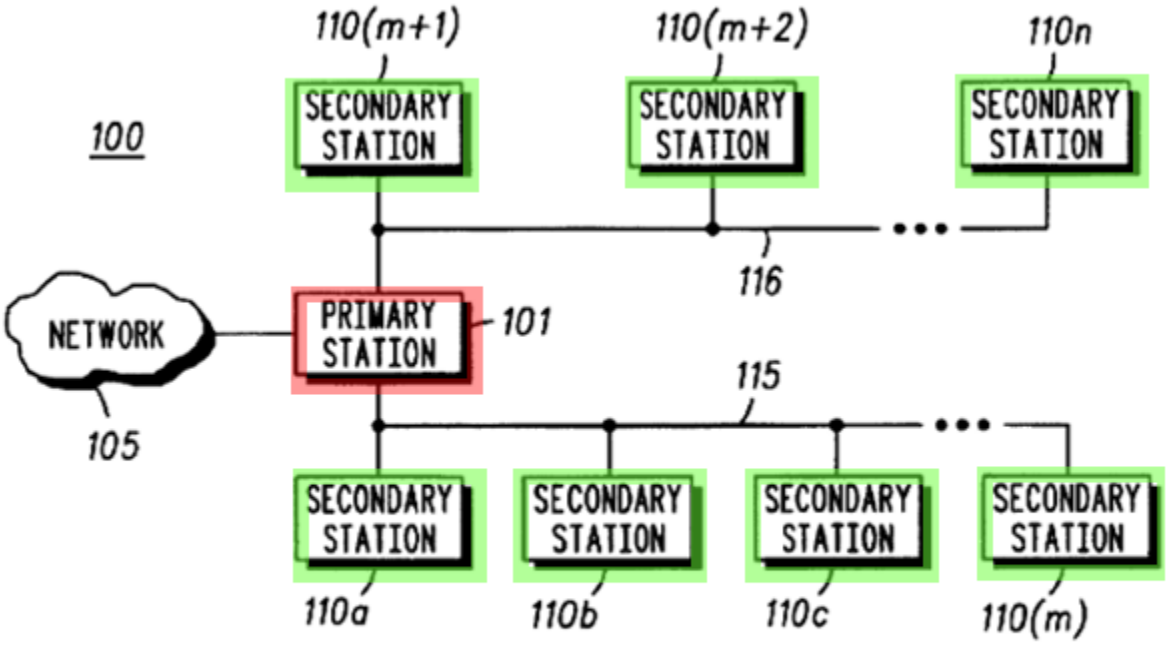


FIG. 1

Ex-1006, Figure 1.

Indeed, Figure 2B of AAPA shows **two masters** connected to a network and to different **slave nodes**. *See* Ex-1003, ¶¶83-85.

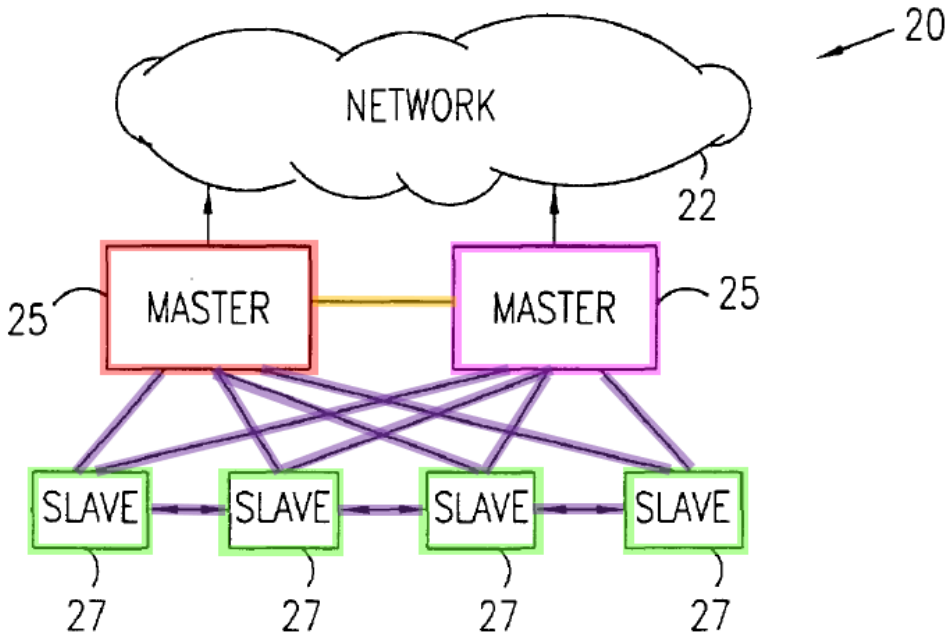


FIG. 2B
PRIOR ART

Ex-1001, Figure 2B.

(1) Reasons To Combine Vink With Patrick (And Further With AAPA)

It would have been obvious to combine the teachings of Vink with Patrick (and AAPA for Ground 2). Vink discloses a communication system that is "known per se," such as "Local Area Networks (LAN's)," Ex-1005, 1:11-15, but Vink is not limited to a LAN. *See* Ex-1005, 22:10-14; Ex-1003, ¶¶74-79.

One well-known exemplary communication system is CableComm™, taught by Patrick, in which the master node can connect to a network, "such as the Internet, online services, telephone and cable networks, and other communication systems.". Ex-1006, 1:19-28, 3:20-24. Ex-1003, ¶80.

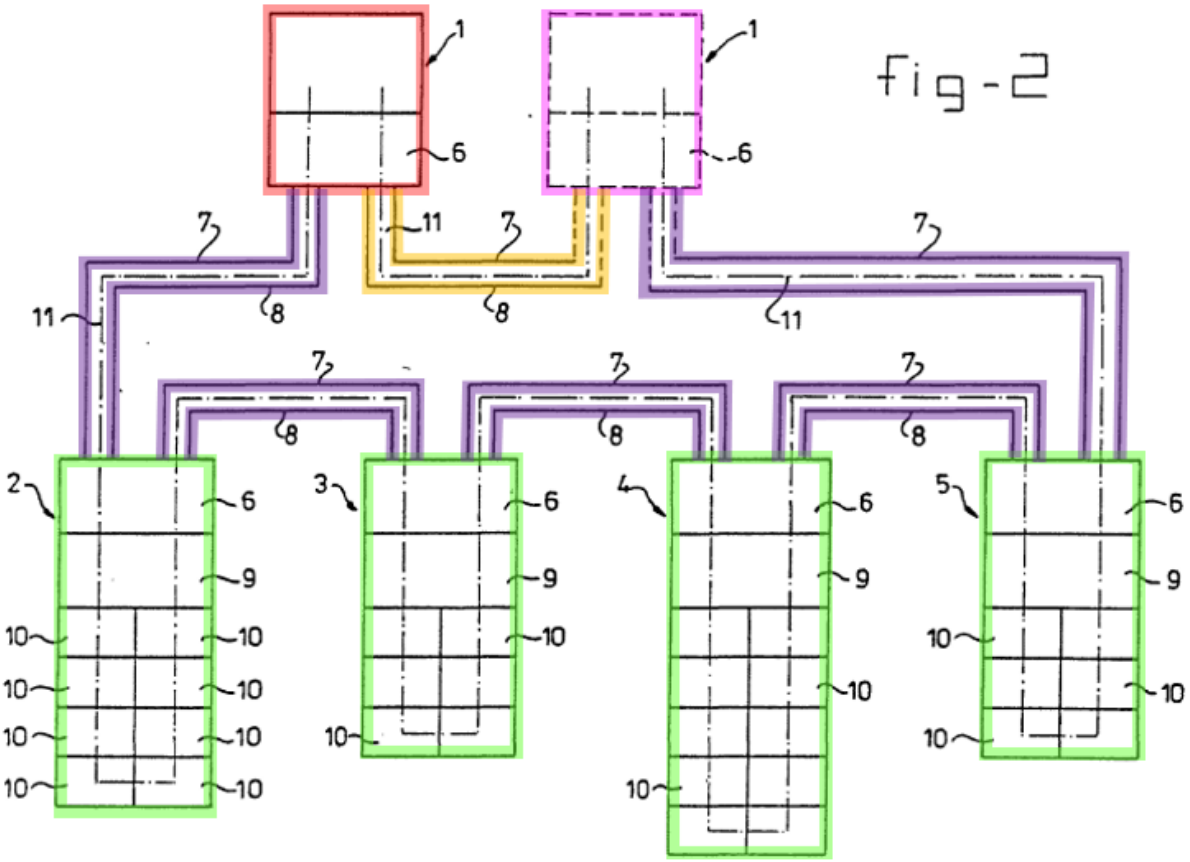
Another admittedly well-known communication system includes Digital Subscriber Line (DSL) which "is the type of infrastructure that links most home and small business subscribers to their telephone service providers." Ex-1001, 1:10-14. As AAPA teaches, "DSL thus opens the most critical bottleneck in local-loop access to high-speed networks, such as Asynchronous Transfer Mode (ATM) and Internet Protocol (IP) networks." *Id.*, 1:14-20. Ex-1003, ¶86.

It would have been obvious to combine the teachings of Patrick (and AAPA for Ground 2) to Vink's communication system by connecting Vink's masters to a known network, such as those taught by Patrick and/or AAPA, to achieve the well-known benefits of gaining access to an external network. *See* Ex-1006, 3:20-24; Ex-1001, 1:14-20. Connecting masters to an external network, such as the Internet (as taught by Patrick) using a specific protocol like the ATM or IP (as taught by AAPA) would have had the benefit of enabling downstream (i.e., data flowing from the external network "down" to the communication system) and upstream (i.e., data flowing "up" from the communication system to the external network) communications. *See* Ex-1006, 1:34-40. And subscribers connected to Vink's communication system would have benefitted by gaining access to a wider array of information that is available on the external network, such as information available on other computers/servers. This was done routinely in the art, and a POSITA would have had a reasonable expectation of success in making such a modification

to Vink's masters given the similarities in the references, including the fact that they are all directed to communication systems that allow master and slave units to communicate with each other. Techniques to further allow the communication system to access external networks by connecting the master to an external network were well-known in the art, as evidenced by Patrick and/or AAPA. Ex-1003, ¶¶74-88.

b) [1a/4a/9a/11a] first [master] and second master units, each comprising a physical interface to a packet-switched network;

Vink with Patrick (and AAPA for Ground 2) suggests [1a], [4a], [9a], and [11a]. Vink discloses **two masters**:



Ex-1005, Figure 2. Specifically, left master 1 is an **active master** and the right master 1 is a **passive master**. *See* Ex-1005, 16:19-24 ("a plurality of masters… illustrated in Figure 2… In such a multi-master concept, one of the masters is in general active in relation to data exchange, while one or more other masters only have a passive role therein."). Ex-1003, ¶¶94-97.

The **two masters** and the **plurality of slave units** exchange data packets via **transmission lines 7, 8**. *See* Ex-1005, 14:28-34 ("at least one master and slaves are designed to transmit or receive packets of data bits via a transmission system ring in one and/or the other direction."), *id.,* 22:1-9, 15:28-35 ("The transmission

lines 7, 8 may, for example, consist of coaxial cable, optical fibre cable and the like."), 1:11-15. Ex-1003, ¶¶95-96.

Vink with Patrick (and AAPA for Ground 2) suggests masters with a physical interface to an external packet-switched network. As explained above in Section VI.B.1.a [1pre/4pre/9pre/11pre], Patrick teaches that the "**primary station 101** is also coupled to a network 105, which may include networks such as the Internet, on line services, telephone and cable networks, and other communication systems." Ex-1006, 3:20-24. *See* Ex-1003, ¶93.

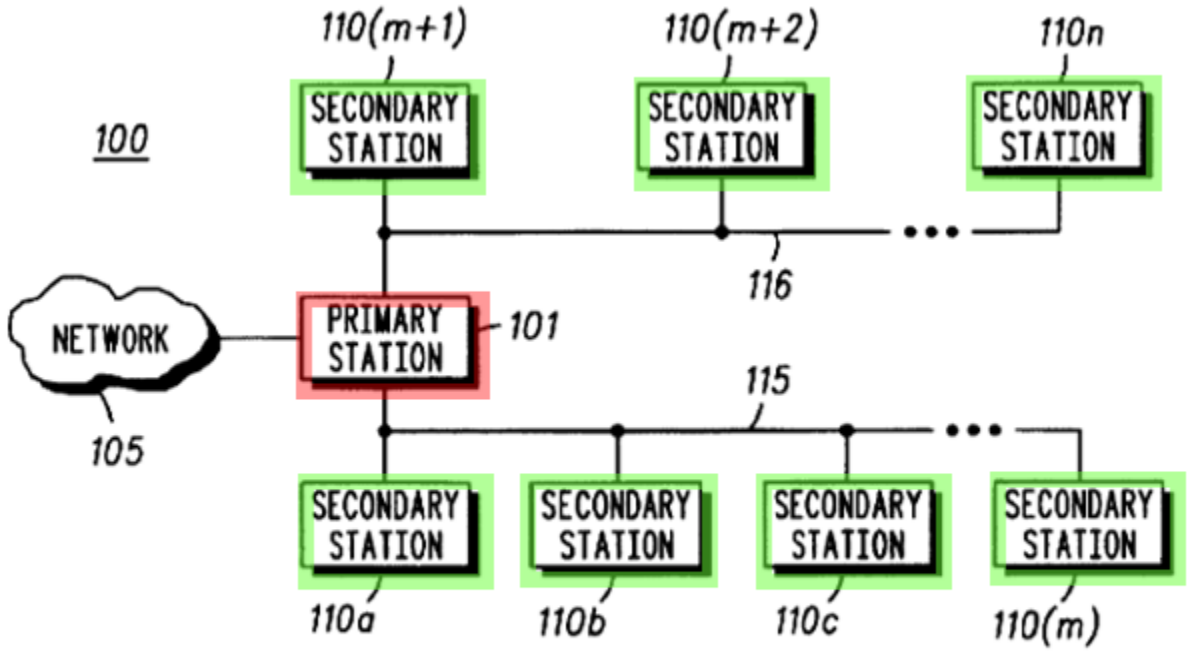


FIG. 1

Ex-1006, Figure 1.

Similarly, AAPA (Figure 2B) shows **two master** nodes connected to a network.[1] Indeed, AAPA teaches that the "**master unit** comprises a core network interface element 24, providing the necessary physical layer (PHY) and data link layer (for example, ATM) functions." Ex-1001, 1:39-42. Ex-1003, ¶¶86-88.

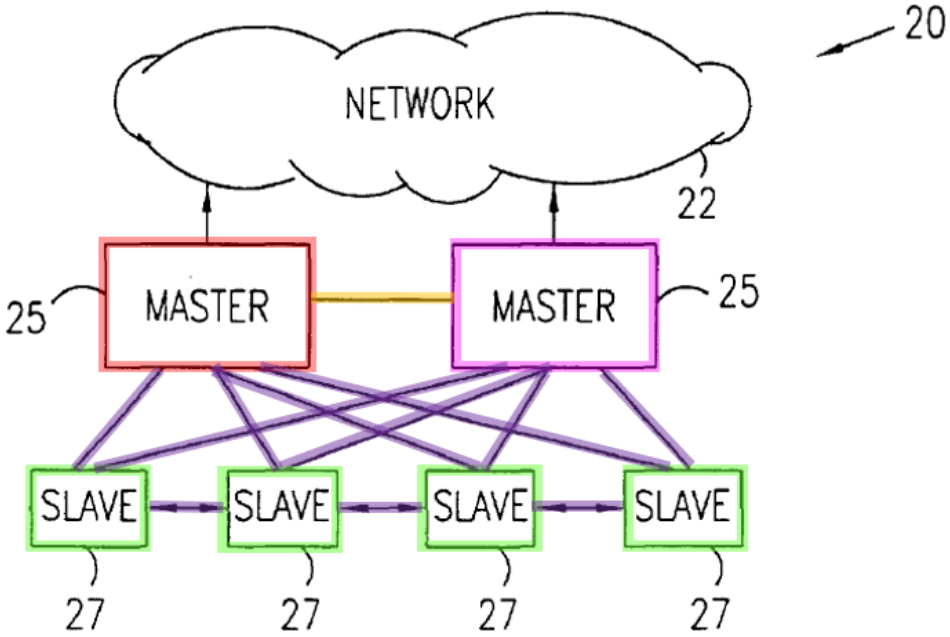


FIG. 2B
PRIOR ART

Ex-1001, Figure 2B.

It would have been obvious to combine the teachings of Vink with Patrick (and AAPA for Ground 2), as more fully discussed in Section VI.B.1.a [1pre, 4pre, 9pre, 11pre].

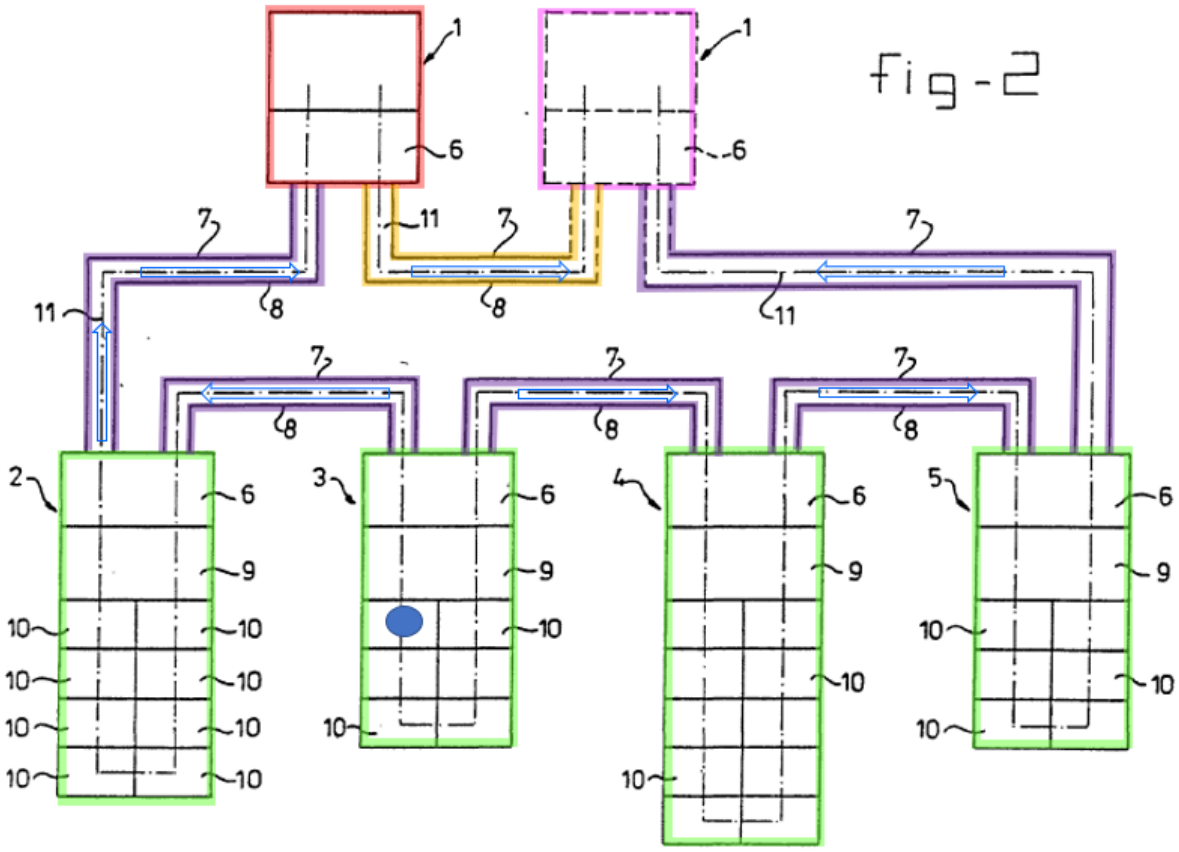
[1] [1a]/[4a]/[9a]/[11a] were admittedly well-known. Indeed, Applicant acquiesced to the Examiner's rejection that AAPA discloses these limitations. *See* Ex. 1002, 159.

Finally, as explained above in Section IV.A [Claim Construction], the term "**second master unit**" should be construed to mean a "standby master that receives upstream packets from the first master unit during normal operations." Vink with Patrick (and AAPA for Ground 2) suggests the claimed "**second master unit**" as properly construed. *See* Ex-1003, ¶¶61-70.

Vink teaches that the "**passive masters** … receive all the information which is intended for the **active master**." Ex-1005, 16:26-31. Such information includes "receiving from the **slaves (2, 3, 4, 5)** via transmission system (7, 8) the data bits originating from the I/O modules (10) [of the slave units]." Ex-1005, 23:22-36. In other words, Vink with Patrick (and AAPA for Ground 2) suggests both **active** and **passive masters** receiving upstream data from the **slave units**. There are only two possible ways the **passive master** can receive all the upstream packets intended for the **active master**. *See* Ex-1003, ¶¶98-102.

Taking upstream packets originating from **slave 3** as an example (see blue circle and blue arrows below), one way to receive upstream packets from **slave 3** would have been to transmit the upstream packets across **slaves 4** and **5** and up to the **passive master 1**. The other way to receive upstream packets from **slave 3** would have been to transmit the upstream packet from **slave 3** to **slave 2**, which passes the upstream packets to **active master 1**, which then sends the upstream packet to **passive master 1** across the **transmission lines 7, 8**. The latter would

have been merely one of two finite ways of having the **passive master 1** "receive all the information which is intended for the **active master**," Ex-1005, 16:26-31. A POSITA would have been able to achieve either transmission of upstream packets given that Vink teaches it is possible to switch the transmission direction. *See, e.g.,* Ex-1005, 14:28-31 ("A very flexible system is obtained … in that the at least one master and slaves are designed to transmit or receive packets of data bits via a transmission system ring in one and/or the other direction."). Indeed, this achieves Vink's stated goal of a "hot stand-by" master which "is of importance … in which failure … may … have serious financial consequences." Ex-1005, 16:34-37. *See* Ex-1003, ¶¶98-102.

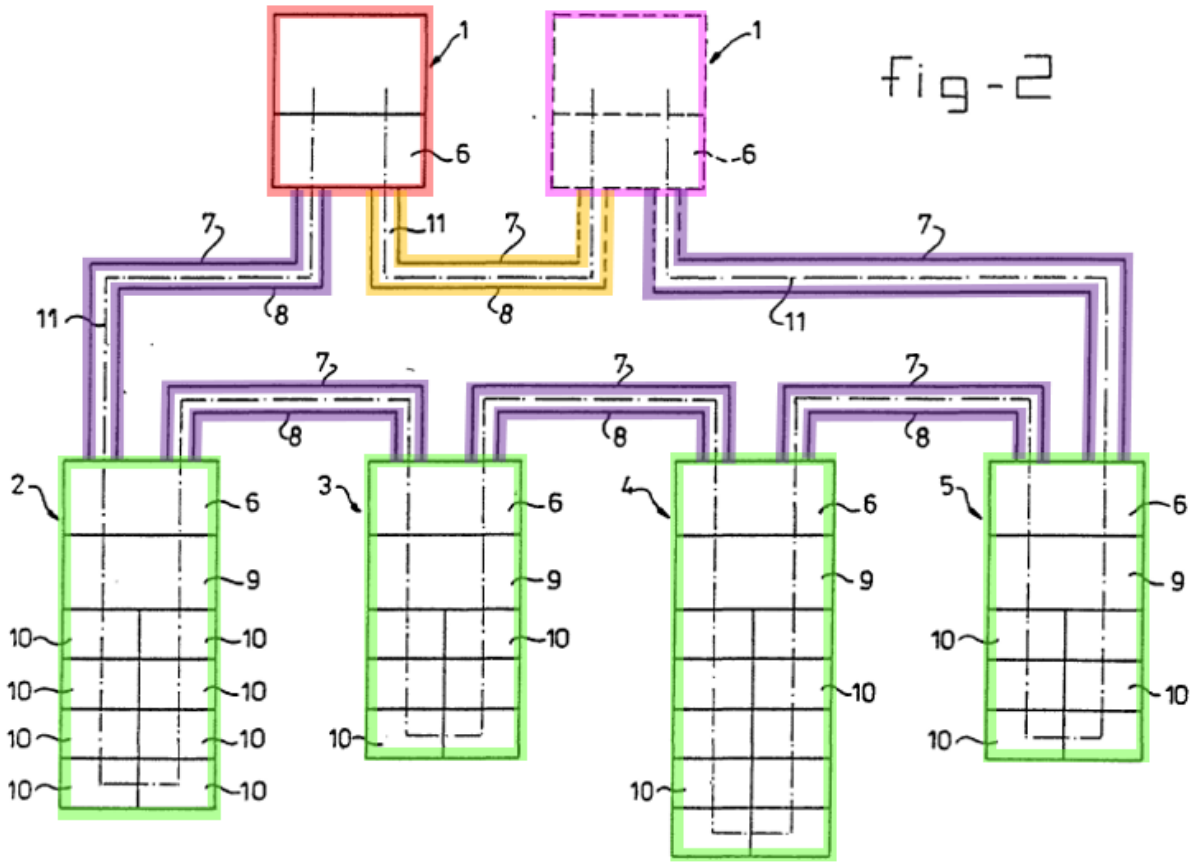


AAPA itself teaches that such "fault protection mechanism" was well-known in the art. The '904 patent states that "[t]he active master bicasts these [upstream] packets, sending them both to network 22 and to standby master 32 over a protection interface 76 between the two masters, as indicated by an arrow 77." Ex-1001, 8:42-47; *id.*, 8:49-52 ("This redundancy in transmission is in accordance with fault protection mechanisms used in high-speed networks ***known in the art***."); *see also* Ex-1009, 52-55. Combining such well-known fault protection mechanism in Vink's communication system, which already has two master stations that has a communication path therebetween, to send upstream

packets during normal operations in an effort to quickly mitigate any fault events would have been a trivial modification that a POSITA would have been motivated to do as well as able to readily implement with a reasonable expectation of success. *See* Ex-1003, ¶102.

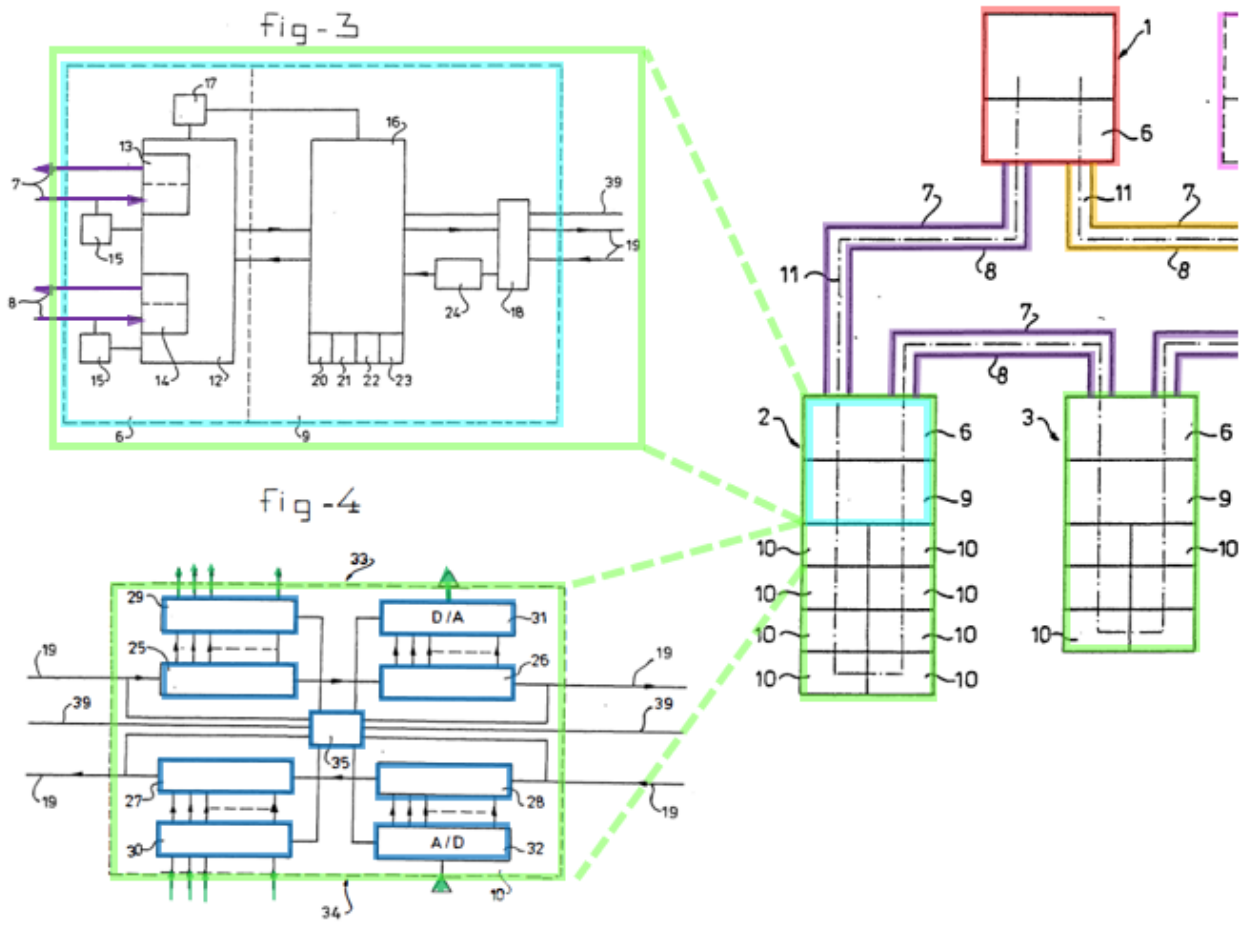
c) [1b/4b/9b/11b] a plurality of slave units, each slave unit comprising one or more ports to respective subscriber lines; and

Vink with Patrick (and AAPA for Ground 2) suggests [1b], [4b], [9b], and [11b]. Vink discloses **slaves 2, 3, 4, and 5**. *See* Ex-1003, ¶¶105-110.



Ex-1005, Figure 2.

Vink teaches that each of the **slaves 2, 3, 4, and 5** includes a **transmitting/receiving unit 6** and **control module 9**, which in turn is connected to one or more **input/output (I/O) modules 10**. Ex-1005, 15:36-16:1.



Ex-1005, Figures 2-4. Specifically, the I/O modules 10 include **inputs (34) and outputs (33)**, such as the parallel set of digital and analog inputs and outputs. *See* Ex-1005, Abstract ("The I/O modules (10) are provided with control means (35)

for recording data bits at the input side (34) and/or presenting data bits at the output side (33).”), *id.,* 19:19-23. These **inputs and outputs** also include connections to **subscriber lines** because the I/O modules 10 are connected to other devices. *Id.,* 7:9-16 (“The sampling instants of I/O modules or ***the devices connected thereto*** can be accurately set by means of the control information supplied, as intended.”). Vink further makes clear that data from the **master 1** flows through the communication system to the **output side 33** and that data from the **input side 34** are supplied to the **master 1**. *See* Ex-1005, claim 1 (“supplying control information to the I/O modules (10) to enable the I/O modules (10): to present in a mutually synchronised manner, at their output side (33), the data bits supplied by the at least one master (1), and to record in a mutually synchronised manner, at their input side (34), data bits for supply to the at least one master (1).”). *See* Ex-1003, ¶105.

Both Patrick and AAPA explicitly teach “**subscriber lines**” and “**ports**.” For example, Patrick teaches that the **secondary stations** can be considered “subscriber access units.” Ex-1006, 1:21-27 (“In the CableComm™ system, a hybrid optical fiber and coaxial cable is utilized to provide substantial bandwidth over existing cable lines to secondary stations such as individual, ***subscriber access units connected to a personal computer, workstation, other data terminal equipment*** (“DTE”), for example, in households having new or preexisting cable

television capability."), *id.,* 1:34-40 ("With the CableComm™ system, digital data may be transmitted … in the downstream direction, from the primary station or controller (connected to a network) to the secondary station of an individual user (***subscriber access unit***."), *id.,* 1:60-64. Indeed, Patrick teaches that the **secondary stations 110a-110n** (i.e., the claimed "slave units") contain an "interface 170, such as an ***ethernet port*** or an RS232 interface, ***for connection to a computer, workstation, or other data terminal equipment***." Ex-1006, 4:59-5:1. Hence, Patrick teaches the claimed "**ports**" (e.g., **ethernet port**) and "subscriber access lines" (e.g., **connections to data terminal equipment**) that are connected to the **slave units** (i.e., the **secondary stations** that provide access to subscribers). *See* Ex-1003, ¶¶106-08.

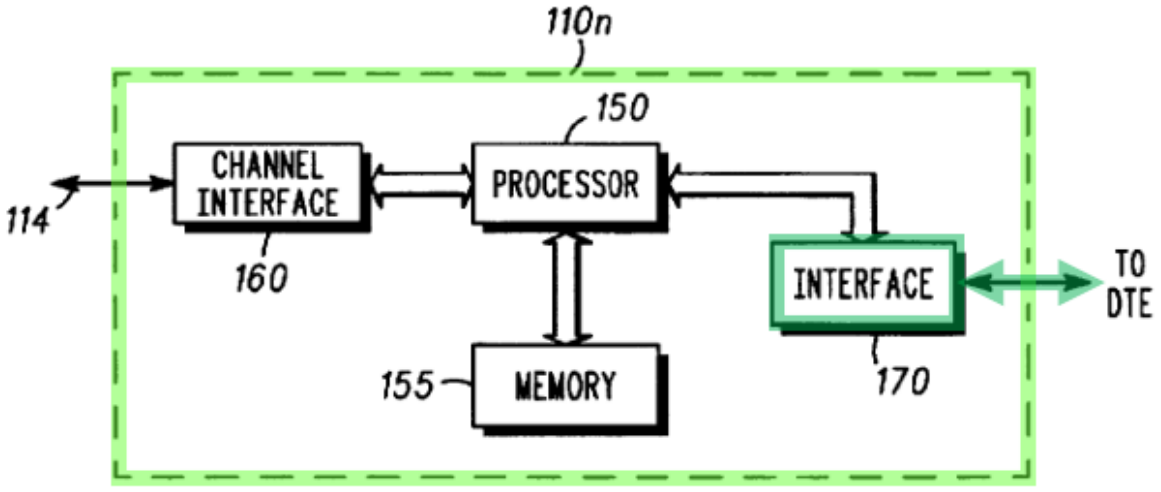


FIG. 3

Ex-1006, Figure 3.

AAPA also teaches [1b], [4b], [9b], and [11b].[2] For example, AAPA teaches that "[e]ach slave unit typically comprises a switching core 29, coupled to a plurality of ports 28 serving respective subscriber premises via suitable DSL modems." Ex-1001, 1:45-48. *See* Ex-1003, ¶109.

It would have been obvious to combine the teachings of Vink with Patrick (and AAPA for Ground 2) to include one or more ports to respective subscriber lines. Providing ports to connect one device (e.g., the slave including an I/O module) to other "devices connected [to I/O modules]" (Ex-1005, 7:15) is one of the most fundamental ways to connect two different devices together. For example, ethernet ports were known long before the priority date of the '904 patent and implementing such ports to the input/output side of Vink's I/O modules in the slave units would have been a trivial modification a POSITA would have been able to implement with a reasonable expectation of success. And having those "devices connected [to I/O modules]" (Ex-1005, 7:15) be connected to subscriber lines (e.g.,

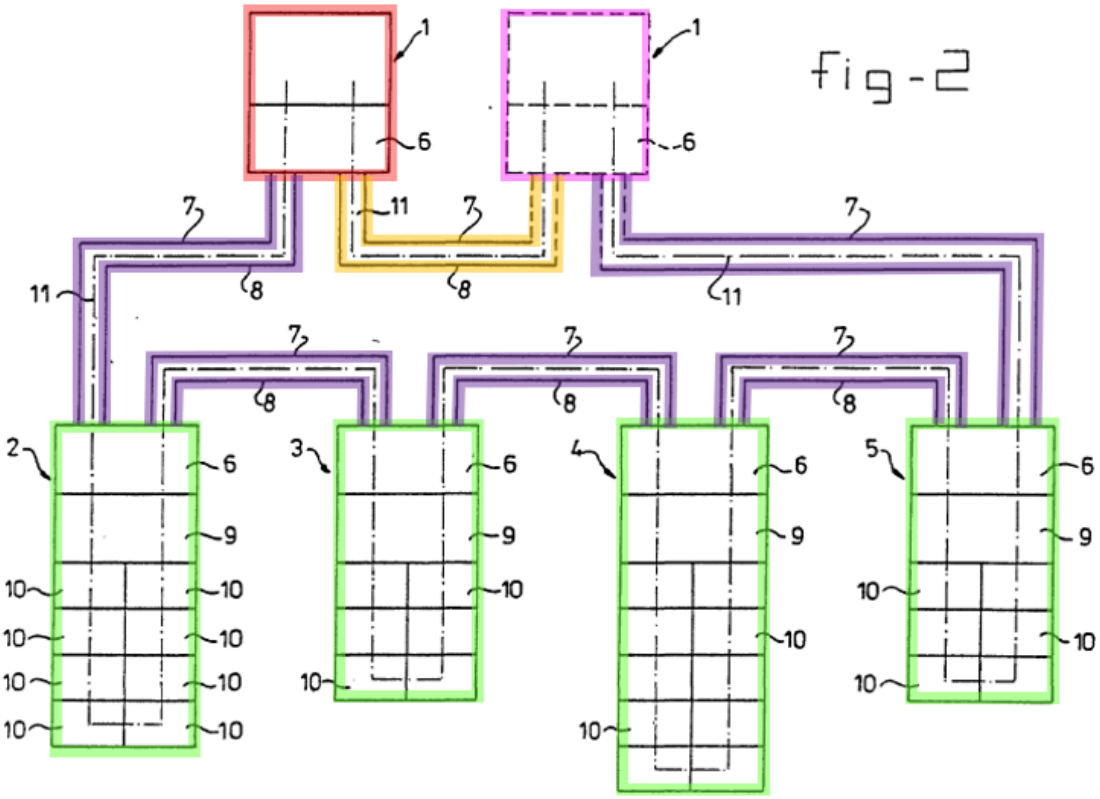
[2] [1b], [4b], [9b], and [11b] were admittedly well-known. Indeed, Applicant acquiesced to the Examiner's rejection that AAPA discloses these limitations. *See* Ex. 1002, 159.

devices found at "subscriber premises," Ex-1002, 159), would have been a trivial modification. *See* Ex-1003, ¶110.

> **d) [1c/4c/9c/11c] a plurality of physical interface lines, which link the slave units in one or more daisy chains, in which the slave units are mutually connected in series by the physical interface lines therebetween, each daisy chain comprising at least a first slave unit connected one of the physical interface lines to the first master unit, […] and a last slave unit connected by another of the physical interface lines to the second master unit.**

Vink discloses [1c], [4c], [9c], and [11c].[3] Vink discloses a plurality of **transmission lines 7, 8** made of coaxial cable, optical fiber cable and the like (i.e., claimed "physical interface lines") which link the **slaves 2, 3, 4, and 5** together. *See* Ex-1005, 15:31-35. *See* Ex-1003, ¶¶112-113.

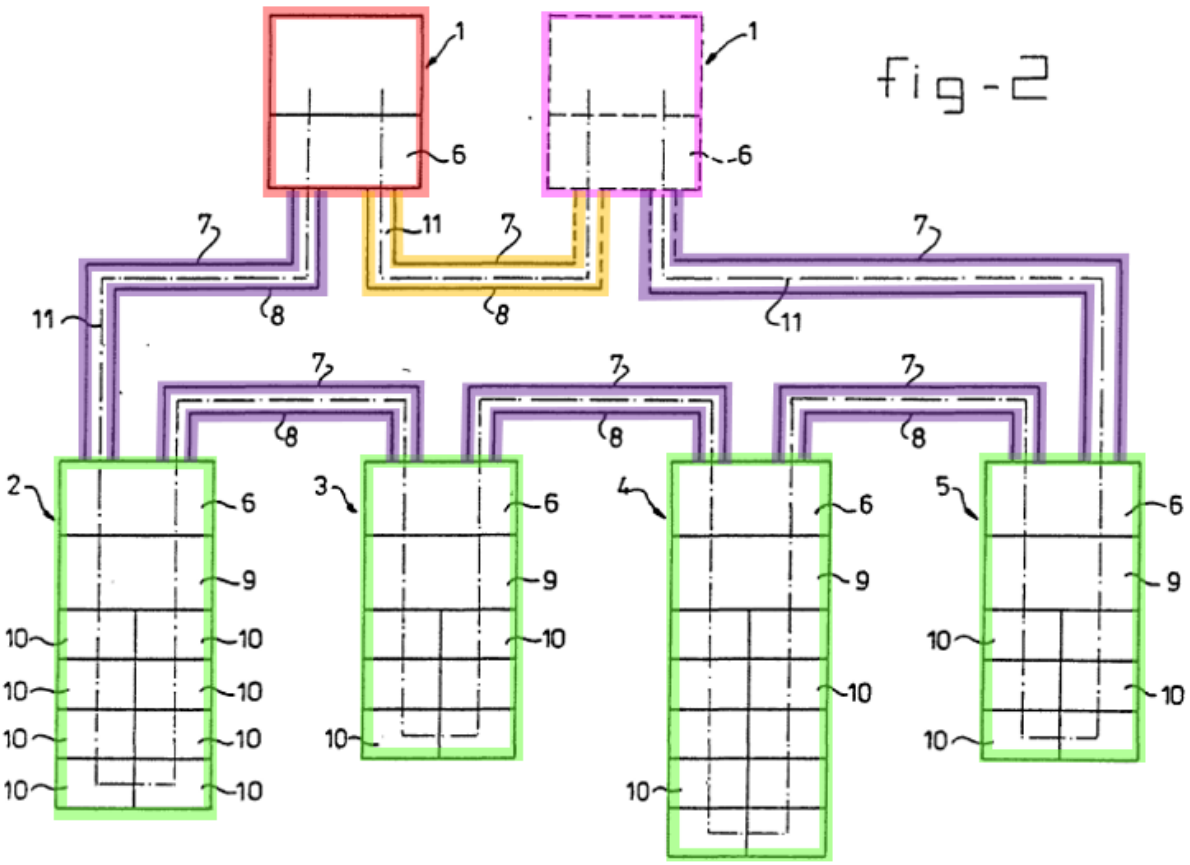
[3] [1c], [4c], [9c], and [11c] were admittedly well-known. Indeed, Applicant acquiesced to the Examiner's rejection AAPA discloses these limitations. Ex. 1002, 159-160.



Ex-1005, Figure 2. As shown, a daisy chain is formed such that the **slaves 2, 3, 4, and 5** are mutually connected in series by the **transmission lines 7, 8** to allow a serial transmission of data. *See id.*, 16:8-10 ("The data bits are serially transferred both in the transmission system 7,8 and to the level of the I/O modules 10."). Specifically, in the daisy chain shown in Figure 2, the **first master 1** is connected to the **first slave 2** via **transmission lines 7, 8**, and the **second master 1** is connected to the **last slave 5** via another set of **transmission lines 7, 8**. *See* Ex-1003, ¶¶113-115.

e) [1d] a second slave unit connected to the first slave unit but not to the first [master unit] or second master unit,

Vink discloses [1d]. As shown below in Figure 2, the **second slave 3** is connected to the **first slave 2**, but not to the **first master** or **second master**.



Ex-1005, Figure 2. *See* Ex-1003, ¶¶117-119.

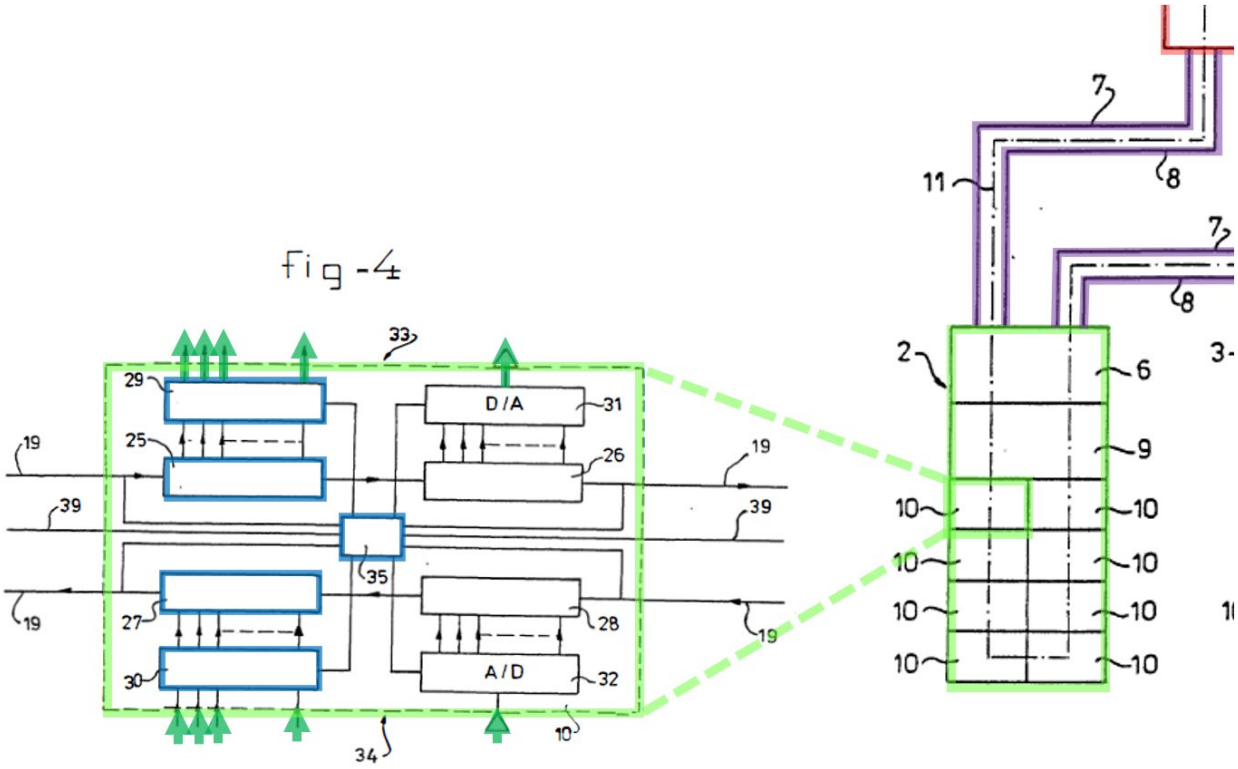
f) [4d] wherein in normal operation, downstream data packets received from the network are passed from the first master unit to each of the daisy chains via the first slave unit in each chain, and upstream data packets received by the slaves in each chain from the subscriber lines are passed

via the first slave unit in the chain to the first master unit for transmission over the network.

Vink with Patrick suggests [4d]. ***First***, Patrick teaches the portion of [4d] reciting "wherein in normal operation, downstream data packets received from the network are passed from the **first master unit** … and upstream data packets received by the **slaves** … are passed … to the **first master unit** for transmission over the network." Patrick teaches that data is normally transmitted in the downstream direction from the network down to the slaves via the master and in the upstream direction from the slave to the network via the master. *See* Ex-1006, 1:34-40 ("With the CableComm™ system, digital data may be transmitted both in the ***downstream direction, from the primary station or controller (connected to a network) to the secondary station*** of an individual user (subscriber access unit), and in the ***upstream direction, from the secondary station to the primary station (and to a network)***."), *id.,* 1:60-64. *See* Ex-1003, ¶¶121-122.

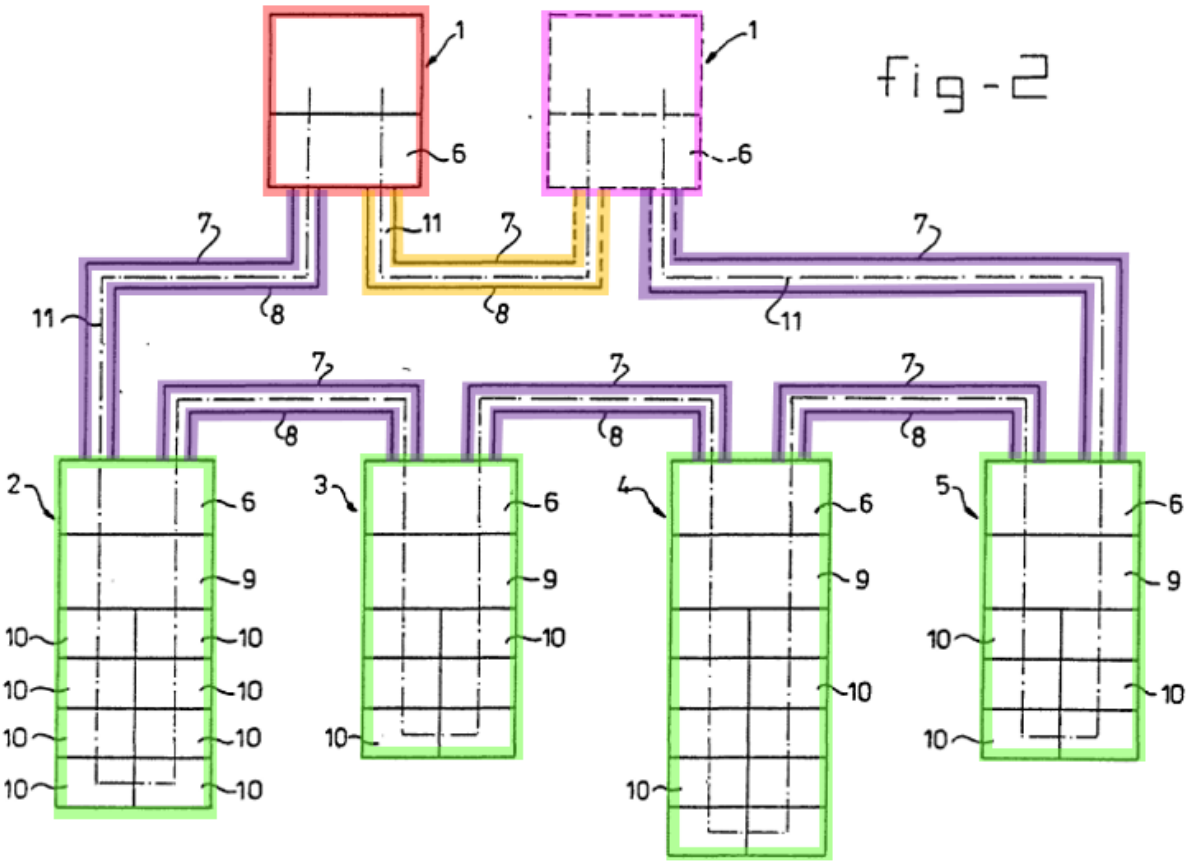
Second, Vink teaches that "data packets … are passed from the **first master unit** to each of the daisy chains via the **first slave unit** in each chain." Vink teaches that data from the **master 1** is supplied to the **output side 33** of the I/O modules 10 (i.e., the ports to the subscriber lines) in each of the **slave units**. *See* Ex-1005, claim 1 ("supplying control information to the I/O modules (10) to enable the I/O modules (10): to present in a mutually synchronised manner, at their

output side (33), the data bits supplied by the at least one master (1), ….”). *See* Ex-1003, ¶121-123.



Hence, a POSITA would have understood that downstream data packets received from the external network (as modified according to the teachings of Patrick) would have been passed from the **first master** into the communication network via the **first slave 2**, shown in Figure 2, to enable a serial bit transfer of the data through the **transmission lines 7, 8** that forms a daisy chain of **slaves 2, 3, 4,** and **5**. *See* Ex-1005, Abstract (“Method and communication system for serially exchanging data… The I/O modules (10) can be connected in groups to the at least

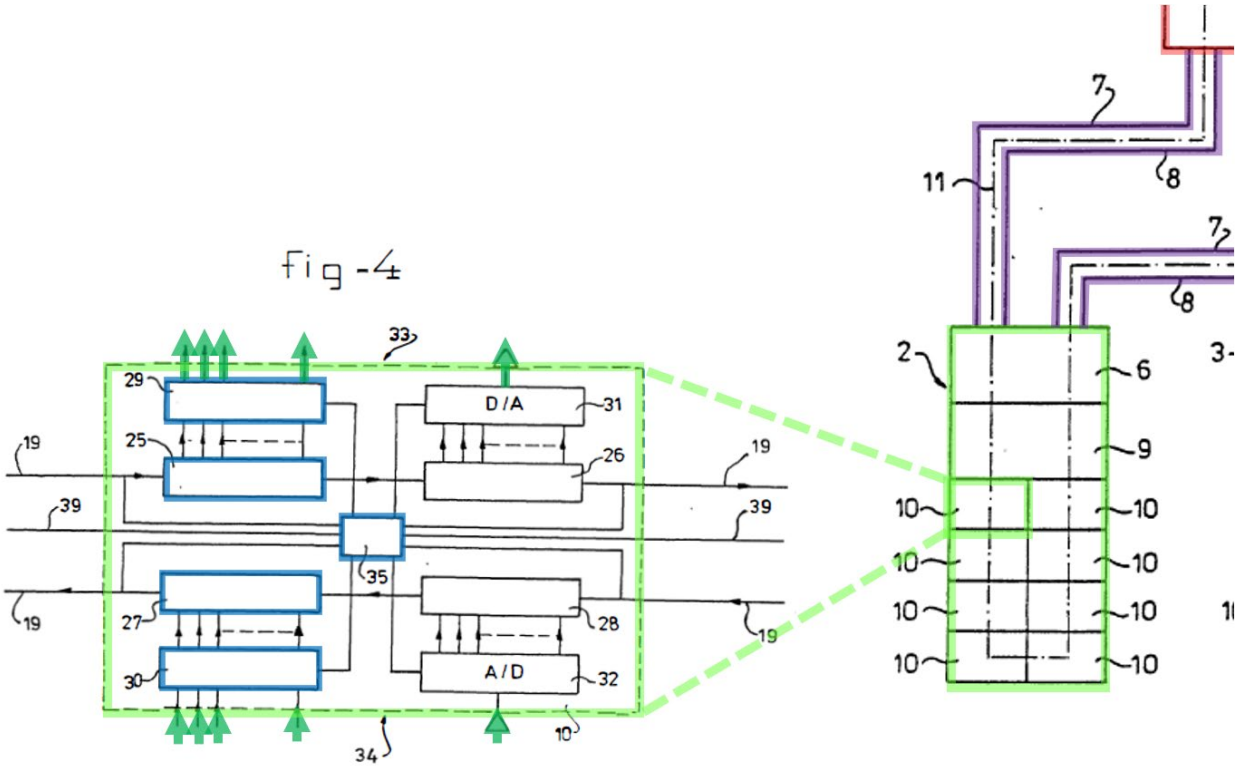
one master (1) via one or more substations (slave) (2, 3, 4, 5) and a transmission system (7, 8) for serial bit transfer."). *See* Ex-1003, ¶¶121-123.



Ex-1005, Figure 2.

Third, Vink teaches that "data packets received by the **slaves** in each chain from the **subscriber lines** are passed via the **first slave unit** in the chain to the **first master unit**." Vink teaches that data from the **input side 34** of the I/O modules 10 (i.e., the ports from the subscriber lines) are supplied to the **master 1**. *See* Ex-1005, claim 1 ("supplying control information to the I/O modules (10) to enable the I/O modules (10): … to record in a mutually synchronised manner, at

their input side (34), data bits for supply to the at least one master (1).”). *See* Ex-1003, ¶121.



Ex-1005, Figure 4. A POSITA would have understood that upstream data packets to be sent to the external network (as modified according to the teachings of Patrick) would have been passed from the **first slave unit 2** to the **first master**. Because Vink’s communication system, which can transmit data in any direction, transmits data serially from the **slaves** to the **master**, upstream data packets to be sent to the external network (as modified according to the teachings of Patrick (and AAPA for Ground 2)) would have been passed from the **first slave unit 2** to the **first master**. *See* Ex-1005, 14:28-31 (discussing that Vink’s communication

system is "designed to transmit or receive packets of data bits via a transmission system ring in one and/or the other direction."), *id.,* 11:9-11 ("The master station and the substations of said communication system each form a node between which information is bit-serial exchanged."), *id.,*16:8-10 ("The data bits are serially transferred both in the transmission system 7, 8 …"). *See* Ex-1003, ¶¶121-123.

Indeed, a POSITA would have readily recognized that it would have been obvious to send downstream data from Vink's **active first master** to the **first slave 2** and upstream data from Vink's **first slave 2** to the **active first master** in Vink's bit serial transmission scheme. It is merely one of a finite number of ways downstream and upstream data can be transmitted across a communication system that utilizes bit-serial transfer of data. And implementing such downstream and upstream data flow would have been an obvious choice with a reasonable expectation of success. *See* Ex-1003, ¶¶121-123.

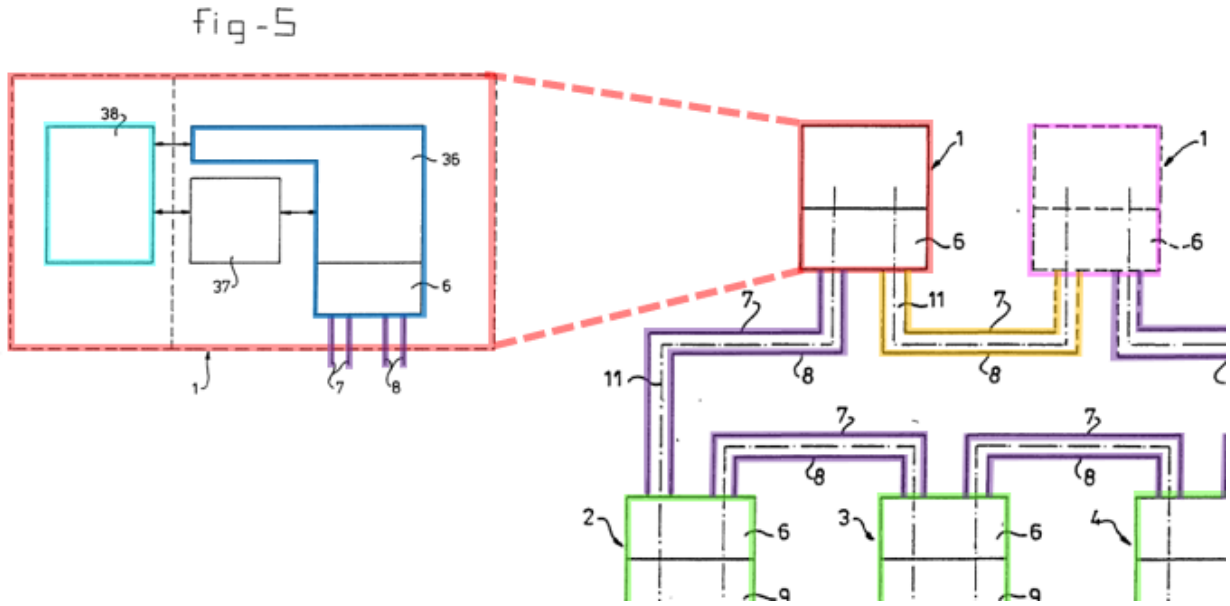
It would have been obvious to combine the teachings of Vink with Patrick to provide downstream and upstream data, in which data flows between an external network and into Vink's communication system through the **master/primary station** and the **slaves/secondary stations** that are connected in series in a daisy chain. Given the daisy-chain configuration taught by Vink in which data is transmitted serially, it would have been obvious to send downstream data from an

external network to the **active first master** through the **first slave 2** during normal operation. Similarly, given the daisy-chain configuration taught by Vink in which data is serially transmitted, it would have been obvious to send upstream data from one of the slave units into the external network by passing the data through **first slave 2** and into the **master** during normal operation. Sending downstream and upstream data in this manner, in which the master and slaves are connected in serial to transmit data serially, would have been a trivial modification that facilitates transport of data into and out of the network. And a POSITA would have been able to implement such traffic flow of data with a reasonable expectation of success. *See* Ex-1003, ¶¶121-123.

g) [9d] wherein each of the first and second master units comprises:

Vink discloses [9d]. Figure 5 of Vink shows the additional components contained in each of the **active master** and **passive master**—namely **transmitting/receiving unit 6**, **control unit 36**, memory 37, and **control computer 38**. *See* Ex-1005, 20:30-38 ("The software for controlling the communication system … is essentially concentrated in the master 1 which comprises, … as shown in Figure 5, a control unit 36 for controlling the data flow via the transmission system 7, 8, with an associated memory 37 for storing and exchanging information with a control computer 38 which contains the

calculations and background information necessary for the process to be controlled via the I/O module."). *See* Ex-1003, ¶125.

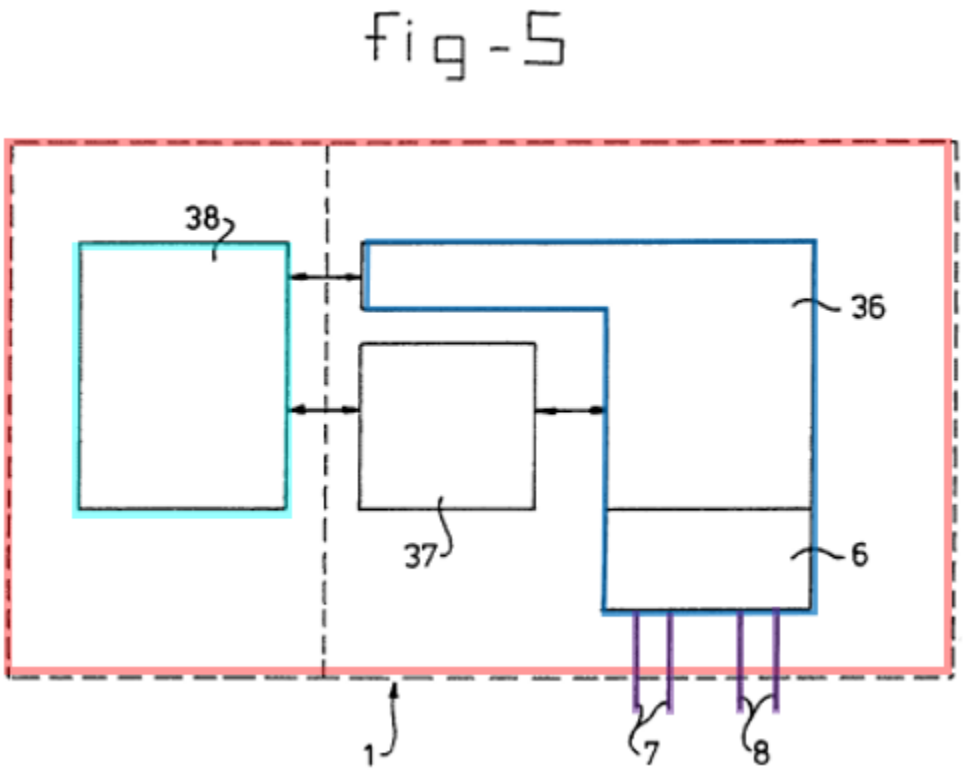


Ex-1005, Figures 2 and 5. Although annotated above to illustrate the components included in only the **active master**, the **passive master** also has the same components. Ex-1005, 16:24-26 ("All the masters are so designed that they can assume the function of the active master, for example in the event of faults."). *See* Ex-1003, ¶125.

h) [9e] a switch[4], configured to route data packets between the respective physical interface and the one or more daisy chains; and

Vink discloses [9e]. Vink discloses that the "**masters**… are capable of functioning as a **switching centre**." Ex-1005, 14:31-33; *see also id.,* 21:10-14. Specifically, the **control unit 36** and **transmitting/receiving unit 6**, located in each of the **active master** and **passive master**, function as the claimed "**switch**." *See* Ex-1003, ¶¶128-130.

[4] The **switch**/**pre-switch** for the **master units** are underlined. The **switch** / **pre-switch** for the **slave units** are ***not*** underlined.



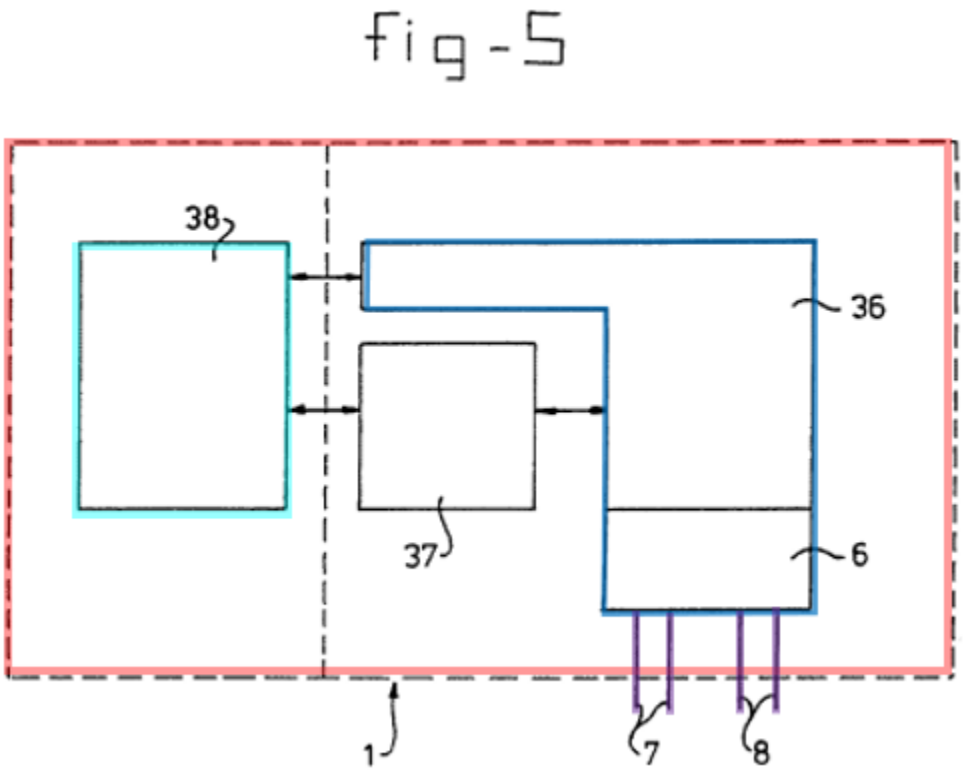
Ex-1005, Figure 5.

The **control unit 36** and **transmitting/receiving unit 6** are "configured to route data packets between the respective physical interface and the one or more daisy chains." For example, Vink explains that the **control unit 36** "extract[s]" information about the particular slave located along the daisy chain the data packets should be sent to and controls the data flow via the **transmission lines 7, 8**. *See* Ex-1005, 21:3-9 ("[T]he control computer 38 provides for and monitors the correct sequence of the data bits for a relevant slave… . The information about this sequence is extracted by the control unit 36 from the status information

received from the relevant slaves."); *id.,* 20:33-34 ("[A] control unit 36 for controlling the data flow via the transmission system 7, 8,"). Moreover, the "**transmitting/receiving unit 6** provides for the monitoring of the transmission" and "the **transmitting/receiving unit 6** also provides for the **switching** of the **transmission ring** or of the communication direction of the data flow." Ex-1005, 21:3-14. *See* Ex-1003, ¶¶128-130.

i) [9f] a pre-switch, which in the event of a fault at a location in one of the daisy chains, re-routes at least a portion of the data packets exchanged with one or more of the slaves in the daisy chain in which the fault has occurred through another one of the daisy chains.

Vink with Patrick suggests [9f]. Vink discloses a **control computer 38**, which functions as a **pre-switch**, in each of the **active master** and **passive master**. *See* Ex-1003, ¶¶131-141.



Ex-1005, Figure 5.

Vink teaches that the control computer 38 "re-routes at least a portion of the data packets exchanged with one or more of the slaves in the daisy chain in which the fault has occurred through another one of the daisy chains." Specifically, Vink explains that "the control computer 38 provides for and monitors the correct sequence of the data bits for a relevant slave, related to the sequence in which the I/O modules are arranged." Ex-1005, 21:4-7. In other words, the control computer 38 determines the routing of data packets to the

correct slave unit based on where the I/O modules are located in the daisy chain. *See* Ex-1003, ¶133.

Vink also discloses that data packets are re-routed in the event of a fault at a location in one of the daisy chains (e.g., switch direction of data flow). *See* Ex-1005, 14:28-34 ("A very flexible system is obtained … in that the at least one master and slaves are designed to transmit or receive packets of data bits via a transmission system ring in one and/or the other direction. Master and/or slaves constructed in this way are capable of functioning as a switching centre so that, for example, the master is able to reach as many connected slaves as possible in a fault situation."); *id.,* 9:4-7 ("It is furthermore possible, in the event of malfunctions in the transmission system, for example during the switching of transmission or communication directions, nevertheless to exchange the correct data bits with the correct slave."); *id.,* 21:10-16. *See* Ex-1003, ¶¶138-140.

Vink further teaches that the communication system can have more transmission lines and slaves than those explicitly shown in Figure 2. *See* Ex-1005, 16:14-17 ("Although a double-ring transmission system is shown, it will be clear that it is possible to work with either a single-ring transmission system or with a transmission system containing more than two rings. Of course, more or fewer slaves can be used."). Although Vink does not explicitly teach two separate daisy chains of slave units, adding a second daisy chain of slave units to that

shown in Figure 2 of Vink would have been a trivial modification. For example, Patrick discloses secondary stations 110a through 110n that are "connected to the primary station 101 on two segments or branches of a communication medium, such as communication media 115 and 116." Ex-1006, 3:24-27. Patrick further teaches that the secondary stations 110a through 110n "may be connected to more than one primary station." Ex-1006, 3:27-29. *See* Ex-1003, ¶¶134-137.

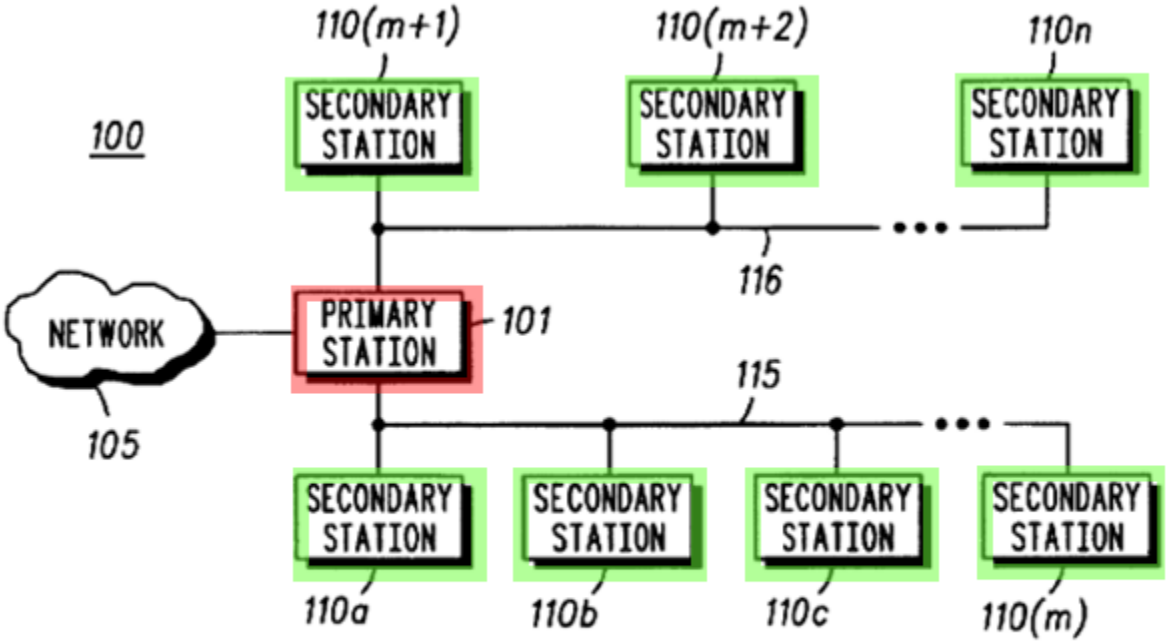
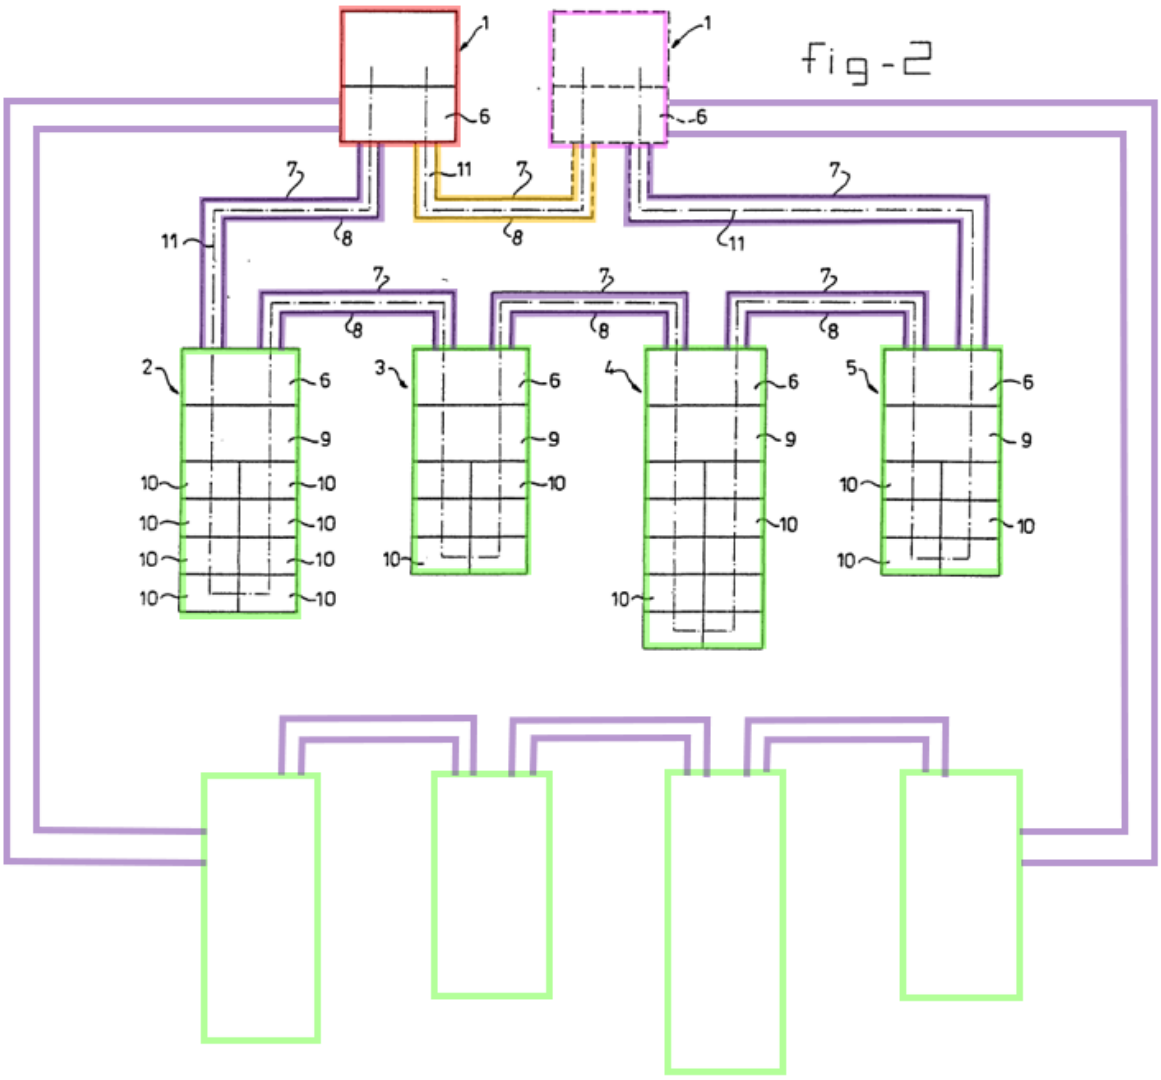


FIG. 1

Ex-1006, Figure 1. Specifically, it would have been obvious to modify Vink to connect additional slaves on additional daisy chains, as schematically illustrated

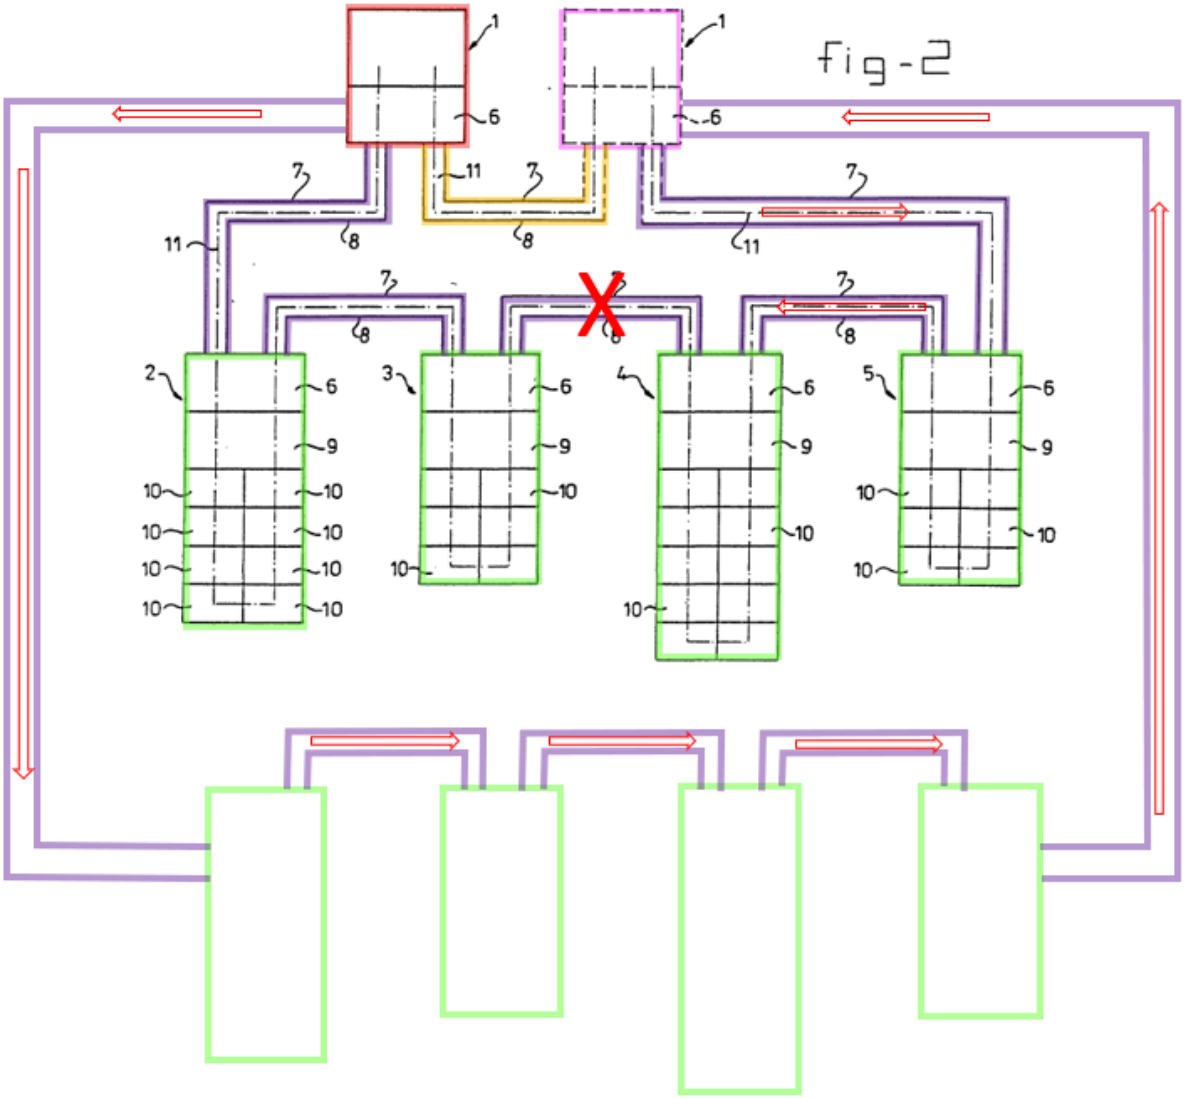
below, in which a second daisy chain of slaves was added to the communication system. For example, POSITA would have been motivated to re-arrange the slaves to decrease the maximum transport delay to slaves located furthest away from the active master. Taking the four slave units shown in Figure 2 as an example, if four more slave units were added to the existing daisy chain, the maximum transport delay for a data packet sent from the active master to the last slave on the daisy chain would be eight hops (i.e., the number of distinct physical interface links the data packet has to traverse). However, if the eight slaves were arranged in two daisy chains, as illustrated in modified Figure 2 below showing four slaves per daisy chain, the maximum transport delay for a data packet sent from the active master to the last slave on each of the daisy chains would be four hops. A POSITA therefore would have understood that such rearrangement of the slaves into multiple daisy chains optimizes and balances the considerations between reducing data packet transport delay and further reduce the possibility of a node failure affecting the chain while simultaneously realizing the benefits associated with the serially connected daisy chain of slaves. *See* Ex-1003, ¶¶134-140.



Ex-1005, Figure 2 (modified). The modification amounts to nothing more than providing additional slaves on one additional "segment[] or branch[] of a communication medium," (Ex-1006, 3:26), to allow connections to additional slaves. A POSITA would have been able to readily modify/program the control computer 38 so that "the correct sequence of the data bits for a relevant slave, related to the sequence in which the I/O modules are arranged," Ex-1005, 21:4-7,

can be sent on either of the transmission lines of the two different daisy chains, such that a data packet can reach the intended slave in the appropriate daisy chain. And because the direction of data flow can be switched in transmission lines 7, 8 as desired, (*see* Ex-1005, 26:4-8 ("Communication system… wherein the at least one master (1) and the slaves (2, 3, 4, 5) are arranged to transmit or receive data bits via a transmission ring (7,8) in one and/or the other direction.")), a POSITA would have been able to program the **control computer 38** to route data to the intended slave across a different daisy chain. This is schematically illustrated below in which "X" denotes a potential fault across the first daisy chain, and the red arrows indicate the direction of downstream data flow, in which the downstream data is routed to slave 4. Indeed, transmission of data across the other daisy chain is merely one out of two possible re-routing of data packets that is possible,[5] and it would at least have been obvious to try to transmit data in either one of the two possible re-routing paths. *See* Ex-1003, ¶¶134-141.

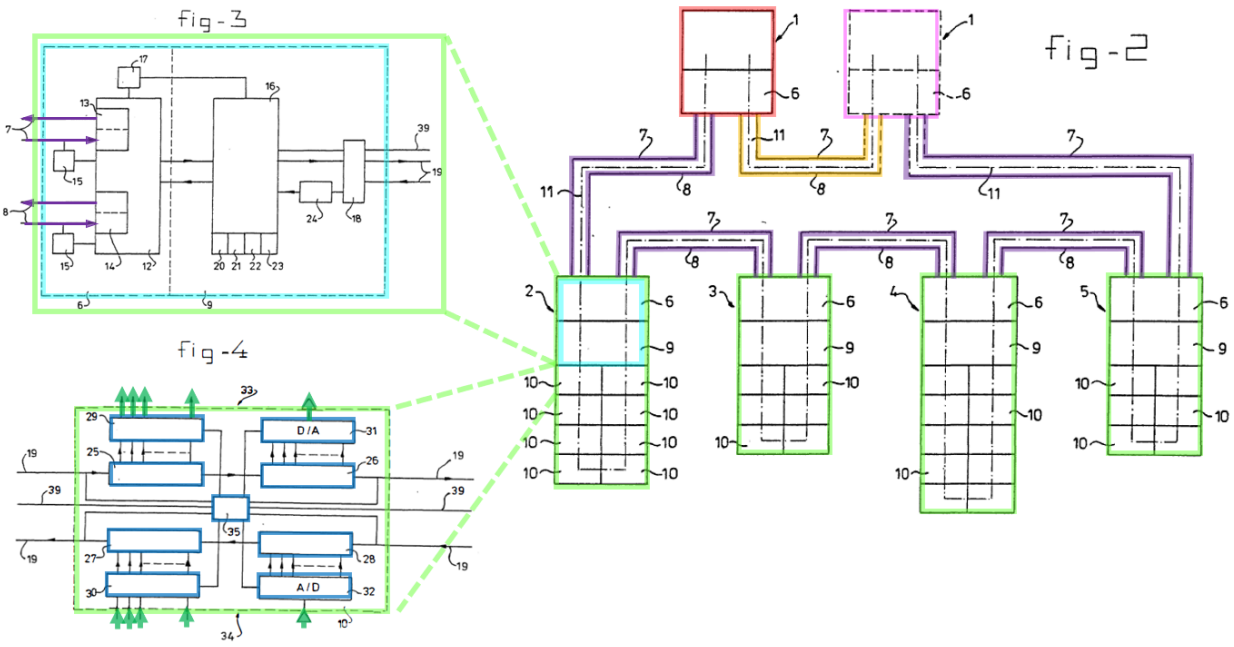
[5] The other possible re-routing involves sending the downstream data from **active master** to **passive master** across **transmission lines 7, 8**. *See* Section VI.B.2.b [5a].



Ex-1005, Figure 2 (modified). And a POSITA would have been able to make such a modification with a reasonable expectation of success because Patrick teaches that "utilizing more or fewer branches, segments or sections of any communication medium," (Ex-1006, 3:30-32), were well within the skill of a POSITA. *See* Ex-1003, ¶135.

j) [11d] wherein each of the slave units comprises:

Vink discloses [11d]. Figures 3 and 4 of Vink show the additional components contained in one of the **slave units**—namely **transmitting/receiving unit 6**, **control module 9**, and **I/O module 10**. *See* Ex-1005, 17:3-4, 15:11-12. *See* Ex-1003, ¶¶142-143.

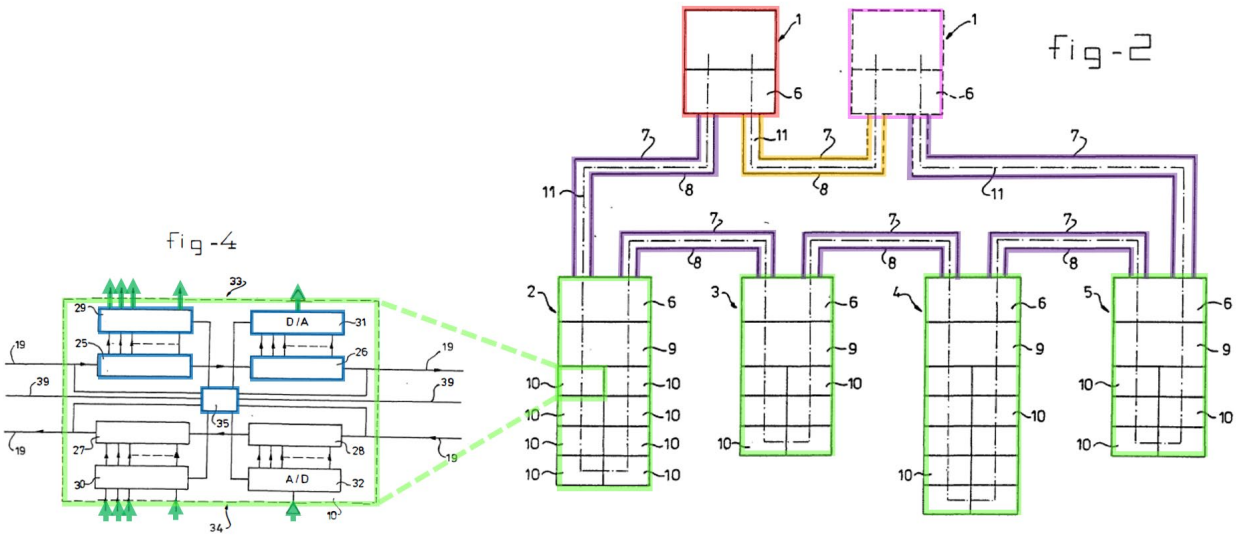


Ex-1005, Figure 2-4.

k) [11e] a switch fabric comprising one or more switches, which convey data packets to respective ports on the switch to which the packets are addressed; and

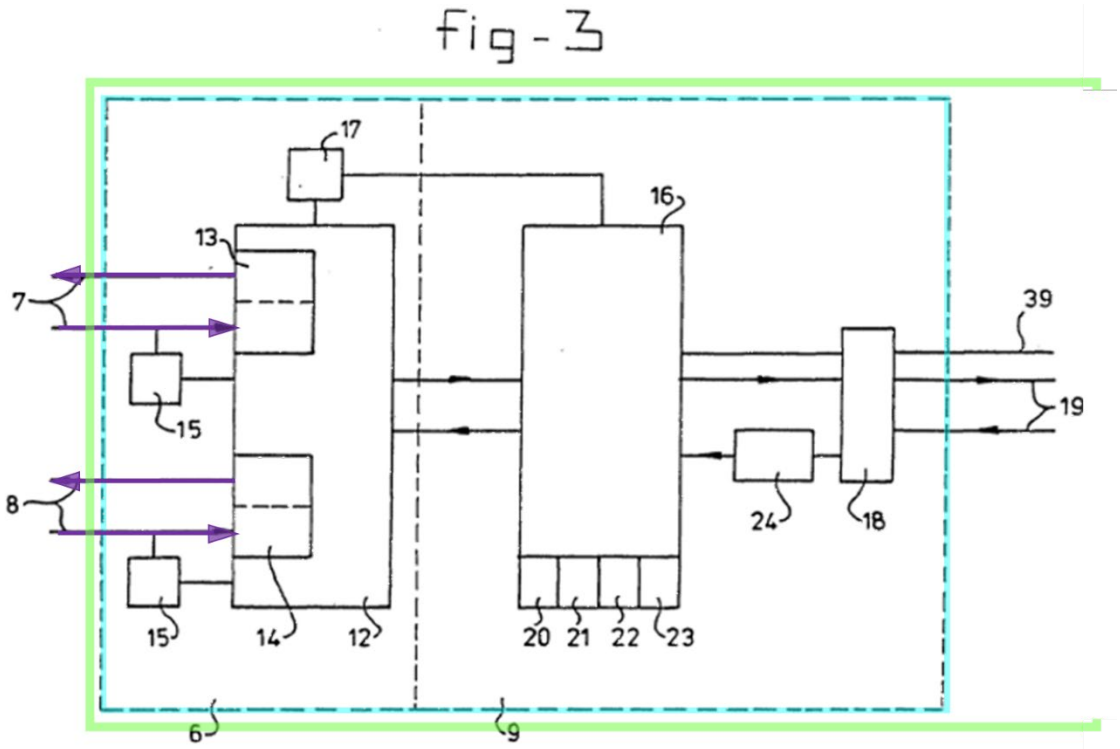
Vink discloses [11e]. Vink discloses that the "***slaves*** … are capable of functioning as a **switching centre**…" Ex-1005, 14:31-33; *see also id.,* 21:10-14. Specifically, **shift registers 25/26**, **buffers 29/30**, and **control means 35** located in

the I/O module 10 of each **slave** together function as the claimed "**switch**" that "convey data packets to respective ports on the switch to which the packets are addressed." *See* Ex-1003, ¶¶145-146.



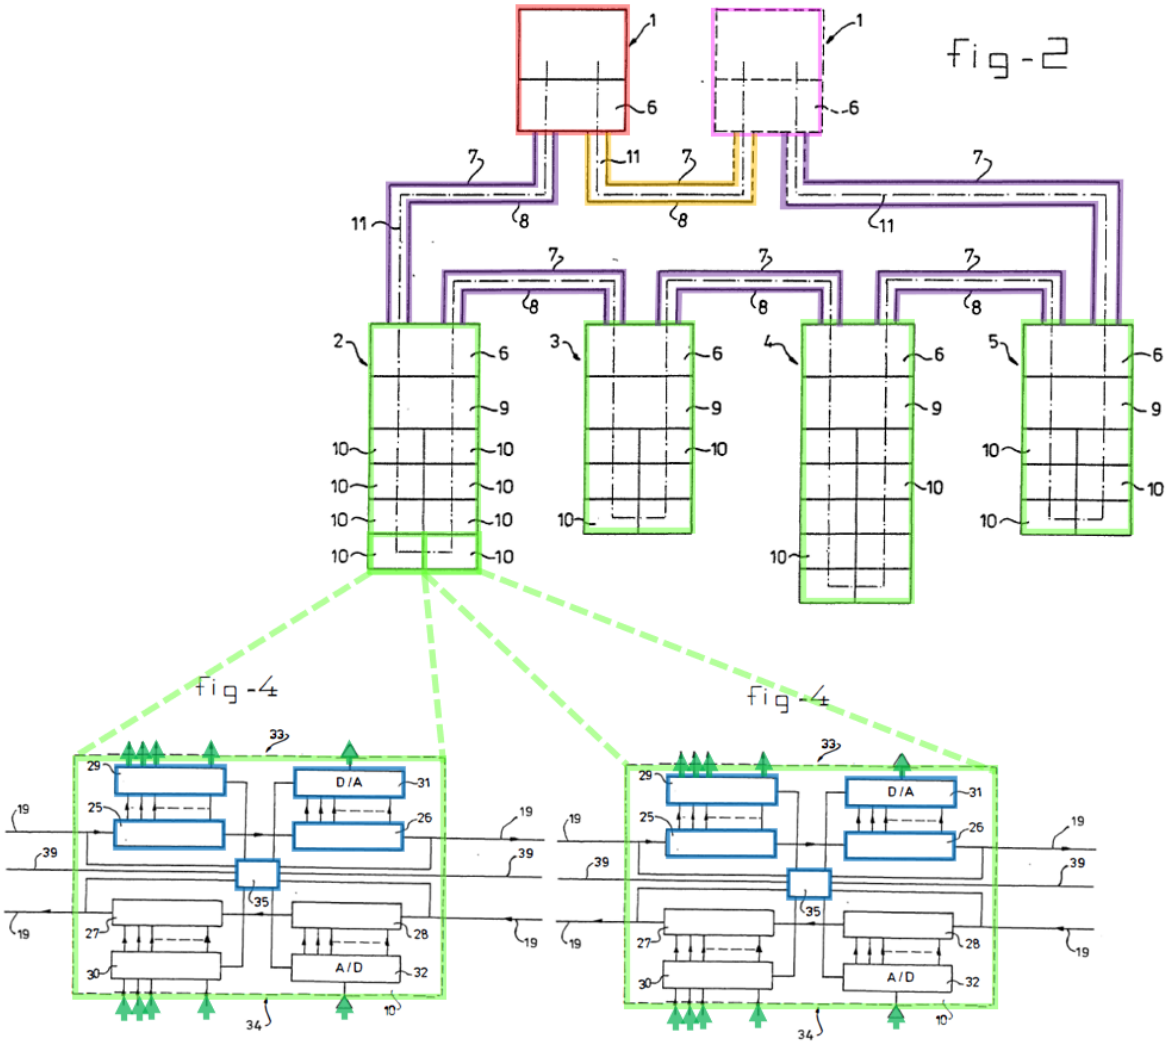
Ex-1005, Figures 2, 4.

To illustrate how these particular components in each I/O module serves as the claimed switches, data flow across the entire slave is illustrated first. As illustrated above, each **slave (2, 3, 4, 5)** includes a **transmitting/receiving unit 6**, **a control module 9**, and one or more **I/O modules 10** (one of which is illustrated to identify the components making up the claimed "**switch**"). As shown below, data packets from **transmission line 7, 8** enter each **slave units (2, 3, 4, 5)** through the **transmitting/receiving unit 6**, which in turn conveys data packet to **control module 9**. *See* Ex-1003, ¶147.



Ex-1005, Figure 3; *see also* Ex-1005, 17:5-8 ("The transmitting/receiving unit 6 comprises a multiplexer 12 consisting of two modems 13, 14 which are respectively connected to the first transmission ring 7 and to the second transmission ring 8."); *id.,* 17:13-15 ("After demodulation by one of the modems 13, 14, the data received are fed to the control module 9 connected to the multiplexer 12 which … is arranged to exchange information with the master…"); *id.,* Figure 3. The **control module 9** then conveys data to each of the **I/O modules 10** via local data bus 19. Ex-1005, 15:36-16:1 ("Each slave 2, 3, 4, 5 comprises a control module 9 which is connected to the associated transmitting/receiving unit 6 and to which one or more input (I) and/or output (O) module(s) are connected.");

id., 17:32-34 ("The respective I/O modules 10 are connected in cascade to a said local data bus 19 (not shown in Figure 3)."). Such a cascade is shown below across two I/O modules in which data is conveyed onto the next **I/O module** via local bus 19. In addition, the **I/O modules** are controlled by control line 39. Ex-1005, 17:34-35 ("A control line 39 is provided for controlling the I/O modules 10."). *See* Ex-1003, ¶¶147-149.



Ex-1005, Figures 2 and 4.

Vink teaches that when data is intended for a particular I/O module 10, the data on local data bus 19 is conveyed to the **output 33** of the I/O module 10. Specifically, within each I/O module 10, **shift registers 25/26** and **buffers 29/30** provide data from the local data bus 19 to the **output 33** of the I/O module 10, under the direction of **control means 35**, e.g., through a control signal line 19 to **buffers 29/30**. Ex-1005, 19:4-8 ("When the data bits intended for the relevant I/O module are all, for example, received in the shift register, the relevant information can be read out, if necessary in parallel, by means of a control signal without disrupting the bit transfer.") ; *id.,* 19:14-18 ("The shift registers 25 and 26 are located in cascade at the output side 33 … Data bits are transmitted from and to the relevant I/O module via the connections, provided with an arrow, to the local data bus 19."). *See* Ex-1003, ¶¶150-152.

Specifically, the signal on control line 39 controls the I/O modules 10 to present data at their **output side 33**. Ex-1005, 17:35-37 ("Control information received is transmitted via said control line 39 to the I/O modules 10 to present them at their output side and/or record data bits at their input side."); *id.,* 19:28-32 ("Connected to the buffers 29, 30 and the converters 31, 32 are control signal lines which are activated from the control module 9 via the control line 39 and the control means 35, respectively, for the presentation of data bits at the output side 33..."); *see also, id.,* 8:25-34. Accordingly, **shift registers 25/26**, **buffers 29/30**,

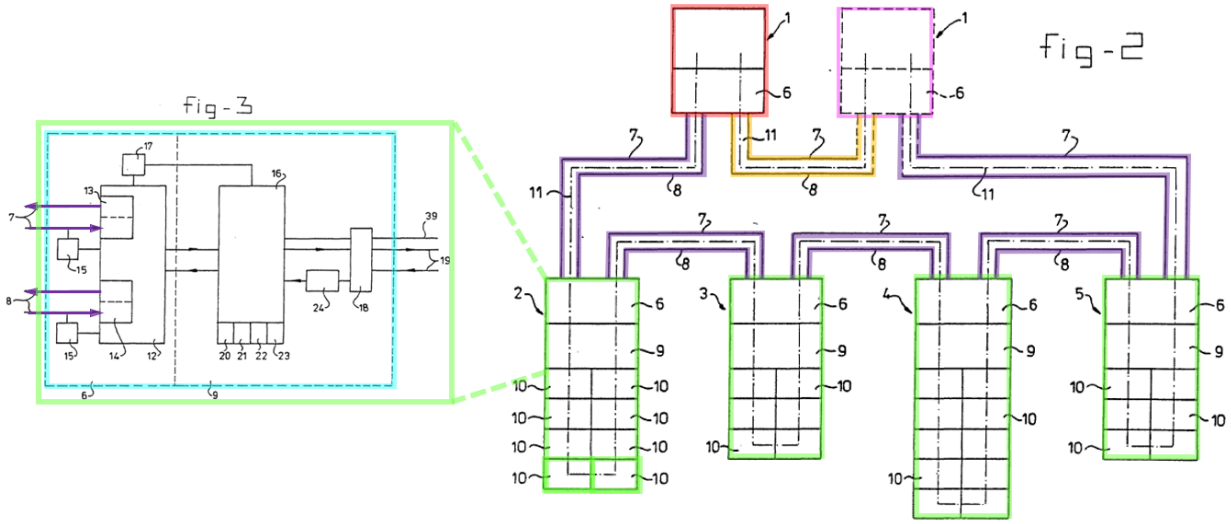
and **control means 35** convey the data packets to the ports on the particular I/O module to which the packets are addressed. Accordingly, Vink discloses the claimed "switch" by the combination of **shift registers 25/26**, **buffers 29/30**, and **control means 35**. *See* Ex-1003, ¶¶150-152.

Moreover, because each I/O module includes its own **shift registers 25/26**, **buffers 29/30**, and **control means 35**, these collection of switches in each I/O module meets the claimed "switch fabric." *See* Ex-1003, ¶152.

> **l) [11f] a pre-switch, which receives the data packets from one of the physical interface lines connected to the slave unit and passes those of the data packets that are addressed to any of the ports on the slave unit to the switch fabric, while passing packets not addressed to any of the ports on the slave unit for output through another of the physical interface lines.**

Vink discloses [11f]. Vink discloses a **transmitting/*receiving* unit 6** and **control module 9** (claimed "**pre-switch**") which receives data packets that are addressed to the I/O modules 10 in the **slave unit** from one of the **transmission lines 7, 8** (claimed "**physical interface lines**"). *See* Ex-1005, 15:36-16:10 ("Each slave 2, 3, 4, 5 comprises a control module 9 which is connected to the associated transmitting/receiving unit 6 and to which one or more input (I) and/or output (O) module(s) are connected. … A dash-dot line 11 diagrammatically shows the data flow in the communication system. The data bits are serially transferred both in the transmission system 7, 8 and to the level of the I/O modules 10."), 25:1-7 ("the

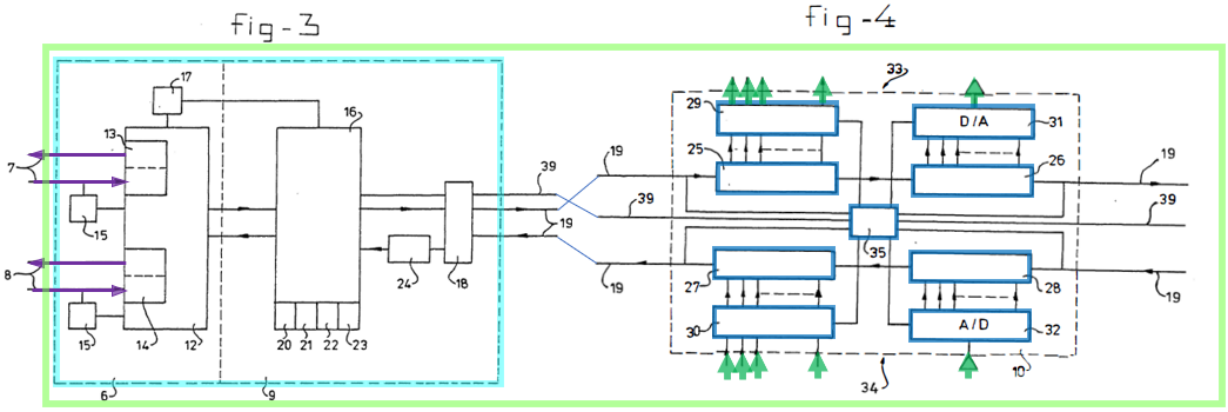
I/O modules (1) being provided with control means (35) and the control module (9) being provided with means (16), coupled to the control means (35), for controlling in a synchronised manner, in response to control information received from the at least one master (1), the inputting and/or outputting of data bits by the connected I/O modules (10).”). *See* Ex-1003, ¶155.



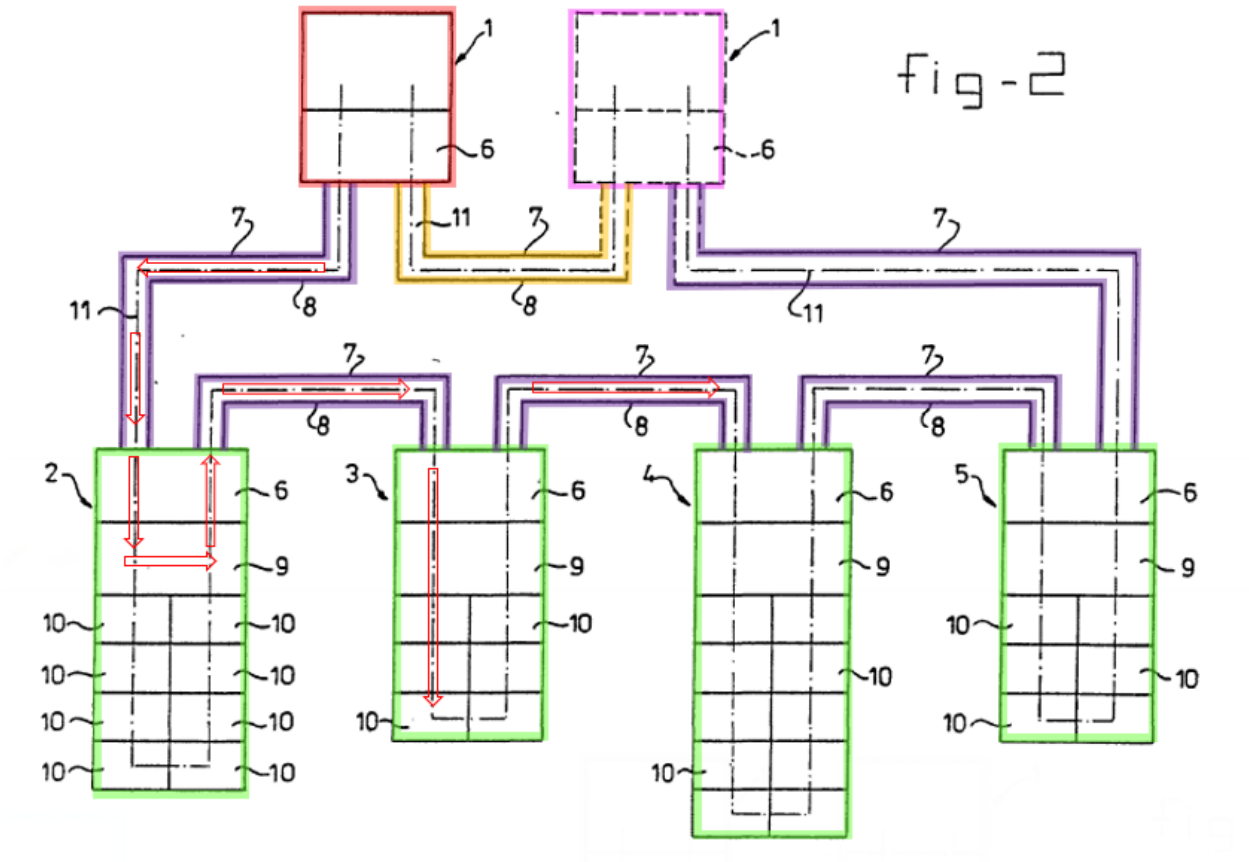
Ex-1005, Figures 2 and 3.

After receiving the data packets, the **transmitting/receiving unit 6** and **control module 9** of a particular **slave** retain only data packets addressed to the **ports (e.g., output 33)** of any I/O module 10 of the particular **slave**, while passing data packets addressed to other **slaves** onto the other **slaves** through the next **transmission lines 7, 8**. *See* Ex-1005, 8:25-34 (“Yet a further embodiment of the method according to the invention, with which **only the data bits intended for**, and originating from[,] the I/O modules connected to **a *local* data bus are**

transferred on said data bus, comprises the steps of: adding address information, under the control of the at least one master (1), to the data bits intended for the I/O modules (10) connected to a respective slave (2, 3, 4, 5), and detecting the address information by the slaves (2, 3, 4, 5) **and only exchanging with the connected I/O modules (10) the data bits intended therefor**."); *id.,* 19:28-32 ("Connected to the buffers 29, 30 and the convertors 31, 32 are control signal lines which are activated from the control module 9 via the control line 39 and the control means 35, respectively, for the presentation of data bits at the output side 33 and the recording of data bits at the input side 34."); *id.,* 25:20-24 ("Communication system … wherein the control module (9) is provided with selection means (20-22) for only transmitting data bits intended for, or originating from the I/O modules (10) connected to said control module (9) via the local data bus (19)."). *See* Ex-1003, ¶¶157-158.



Ex-1005, Figures 3 and 4. In other words, Vink discloses that only the data bits intended for a particular I/O module in one **slave** (e.g., **slave 2**) are transferred on local data bus 19/39 to be directed to, for example, **output 33** of that particular I/O module 10 through the **shift registers 25/26**, **buffers 29/30**, and **control means 35**, whereas the data bits intended for other **slaves** (e.g., **slave 3**) are passed on to the **transmission lines 7, 8** between **slaves 2** and **3**. The figure below illustrates how data (in red arrows) intended for one of the **I/O modules 10** in **slave 3** passes through **slave 2** through the **transmitting/receiving unit 6** and **control module 9** of **slave 2**, without passing through the **I/O modules 10** of **slave 2**. *See* Ex-1003, ¶156.



Ex-1005, Figure 2.

2. Dependent Claim 5

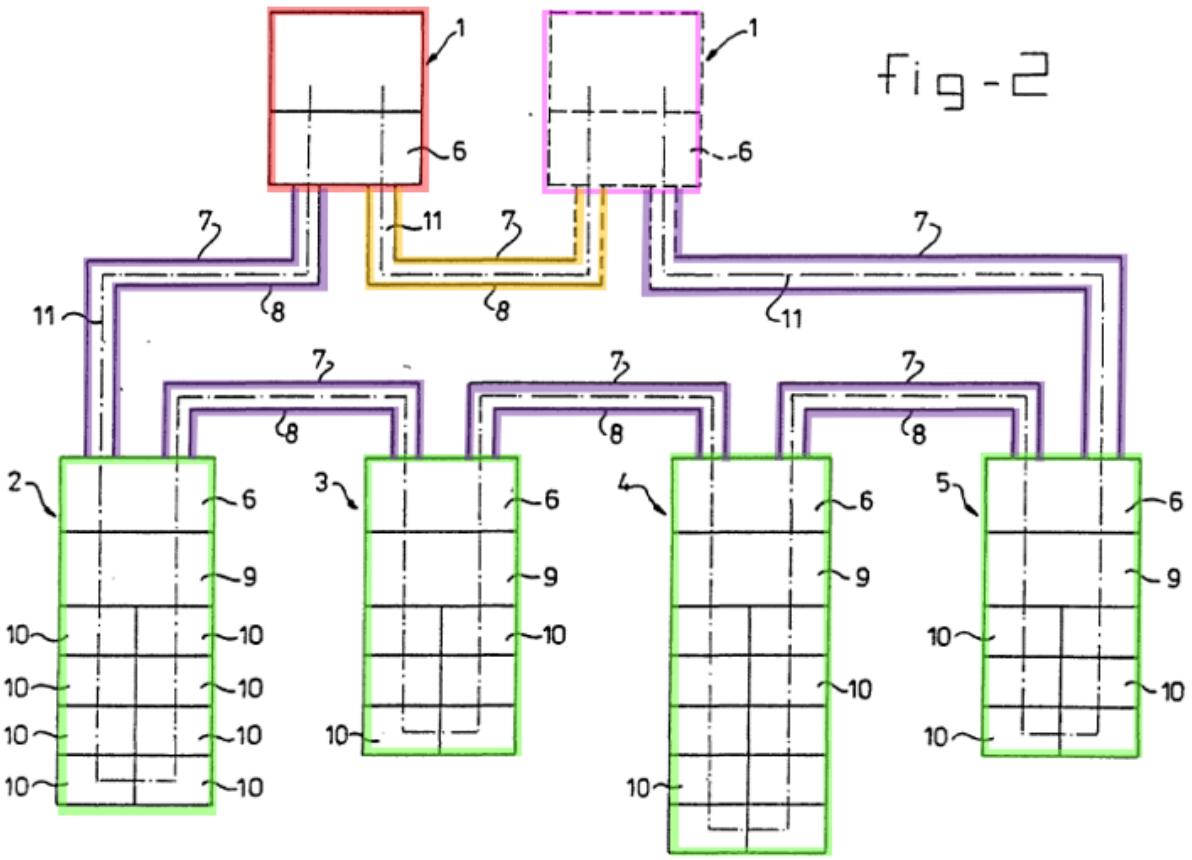
a) [5pre] Apparatus according to claim 4, and

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 4. *See* Section VI.B.1.

b) [5a] comprising a protection interface, which couples the second master unit to the first master unit, and over

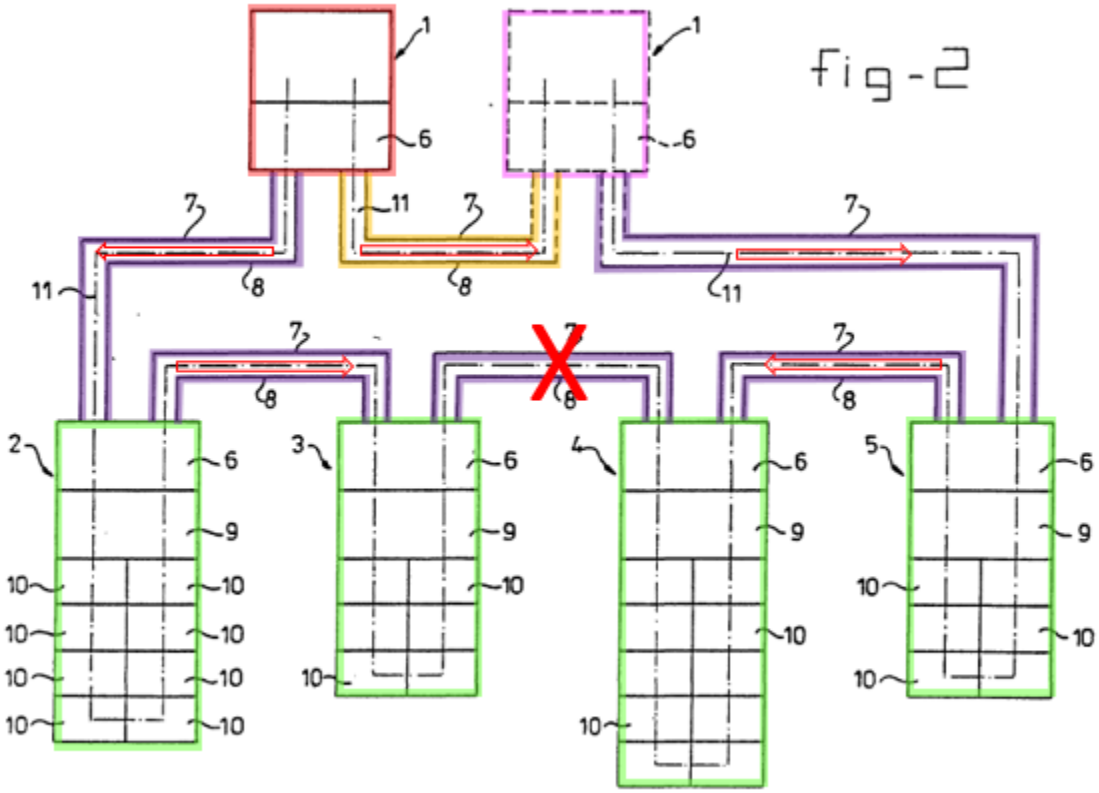
which interface data packets are conveyed between the first and second master units in case of a fault.

Vink alone or Vink with AAPA suggests [5a]. As shown below, Figure 2 of Vink shows **transmission lines 7, 8** that couple the **active master 1** and **passive master 1** together. *See* Ex-1003, ¶¶160-161.



Ex-1005, Figure 2. Vink teaches that the **passive master 1** is designed such that it can assume the role of the **active master 1** when a fault occurs by receiving all the data that the **active master 1** receives. Ex-1005, 16:21-31 ("In such a multi-master concept, one of the masters is in general active in relation to data exchange, while

one or more other masters only have a passive role therein. All the masters are so designed that they can assume the function of the **active master**, for example ***in the event of faults***. The **passive masters** are easily able to detect errors of the active master because, as a consequence of their cascade arrangement and the serial transfer of data bits in the communication system …, ***they receive all the information which is intended for the active master***."); *id.*, 24:24-34 ("Method … comprising under the control of the at least one master (1): forming a packet of data bits … via the **transmission system (7,8)** between the at least one master (1) … information can be exchanged, between packets of data bits and/or control information."). As such, to the extent not explicitly taught, Vink suggests data packets are conveyed between the **active master** and the **passive master** over the **transmission lines 7, 8** in the event of a fault. Such a fault event is schematically illustrated below, with the red arrows depicting the direction of downstream data flow and "X" denoting a fault in **transmission lines 7, 8** between **slaves 3** and **4**. *See* Ex-1003, ¶¶161-164.



Ex-1005, Figure 2. As shown, to send data packets intended for **slaves 4** and **5**, it would have been obvious to send data between **active master** and the **passive master** over **transmission lines 7, 8** to allow transmission of data to all of the **slaves (2, 3, 4, 5)** in the daisy chain. Having recognized that a fault exists in the **transmission lines 7, 8** between **slaves 3** and **4**, one way to ensure that the **passive master** "receive[s] all the information which is intended for the active master" so that the **passive master** "can assume the function of the active master," (Ex-1005, 16:24-31), is to transmit the data packets across the **transmission lines 7, 8** located

between **active master** and the **passive master**. This would have merely been one of a finite number of ways to transmit data between the **active master** and the **passive master** and a POSITA would have had a reasonable expectation of success in achieving such a delivery of data based on the teachings of Vink. *See* Ex-1003, ¶¶161-164.

To the extent Vink does not explicitly teach [5a], Vink with AAPA suggests this limitation. The '904 patent states that "[t]he active master bicasts these packets, sending them both to network 22 and to standby master 32 over a protection interface 76 between the two masters, as indicated by an arrow 77. The standby master also transmits the upstream packets over network 22, as indicated by an arrow 78." Ex-1001, 8:42-49. The '904 patent admits that this is a well-known feature that is commonly used in fault protection mechanisms. *See id.*, 8:49-52 ("This redundancy in transmission is in accordance with fault protection mechanisms used in high-speed networks known in the art, such as the standard '1+1 APS' (automatic protection switching) technique used in SONET."); *see also* Ex-1009, 52-55. A POSITA would have been motivated to implement such 1+1 APS fault protection mechanism between the two masters and between the two masters and the network in case one of the links between the masters or the links from the masters to the network fails. Combining such well-known fault protection mechanism in Vink's communication system, which already has two

master stations that has a communication path therebetween, and as modified with Patrick/AAPA, which connects the master units to a network, would have been a trivial modification that a POSITA would have been motivated to do as well as able to readily implement with a reasonable expectation of success. *See* Ex-1003, ¶165.

3. Dependent Claim 6

a) [6pre] Apparatus according to claim 5,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 5. *See* Section VI.B.2.

b) [6a] wherein the first master unit bicasts the upstream data packets that it receives from the slave units to the network and, via the protection interface, to the second master unit, which transmits the upstream data packets to the network.

As explained in Section VI.B.1.f [4d], Vink with Patrick suggests sending "upstream data packets that it receives from the slave units to the network."

As explained in Section VI.B.2.b [5a], Vink further suggests that "the first master unit []casts the upstream data packets that it receives from the slave units, via the protection interface, to the second master unit."

Finally, Vink with Patrick suggests that the passive master transmits the upstream data packets to the network because "[a]ll the masters are so designed that they can assume the function of the active master, for example in the event of

faults." Ex-1005, 16:24-26. Hence, because the **active master** is designed to transmit upstream data (as modified with Patrick) to the network, a POSITA would also have designed the **passive master** to transmit upstream data packets to the network. *See* Ex-1003, ¶168.

Indeed, as explained in Section VI.B.2.b [5a], AAPA teaches [6a] because the '904 patent admits that this is a well-known feature that is commonly used in fault protection mechanisms. And Vink with Patrick (and AAPA) suggests [6a]. *See* Ex-1003, ¶¶169-170.

4. Dependent Claim 7

a) [7pre] Apparatus according to claim 4,

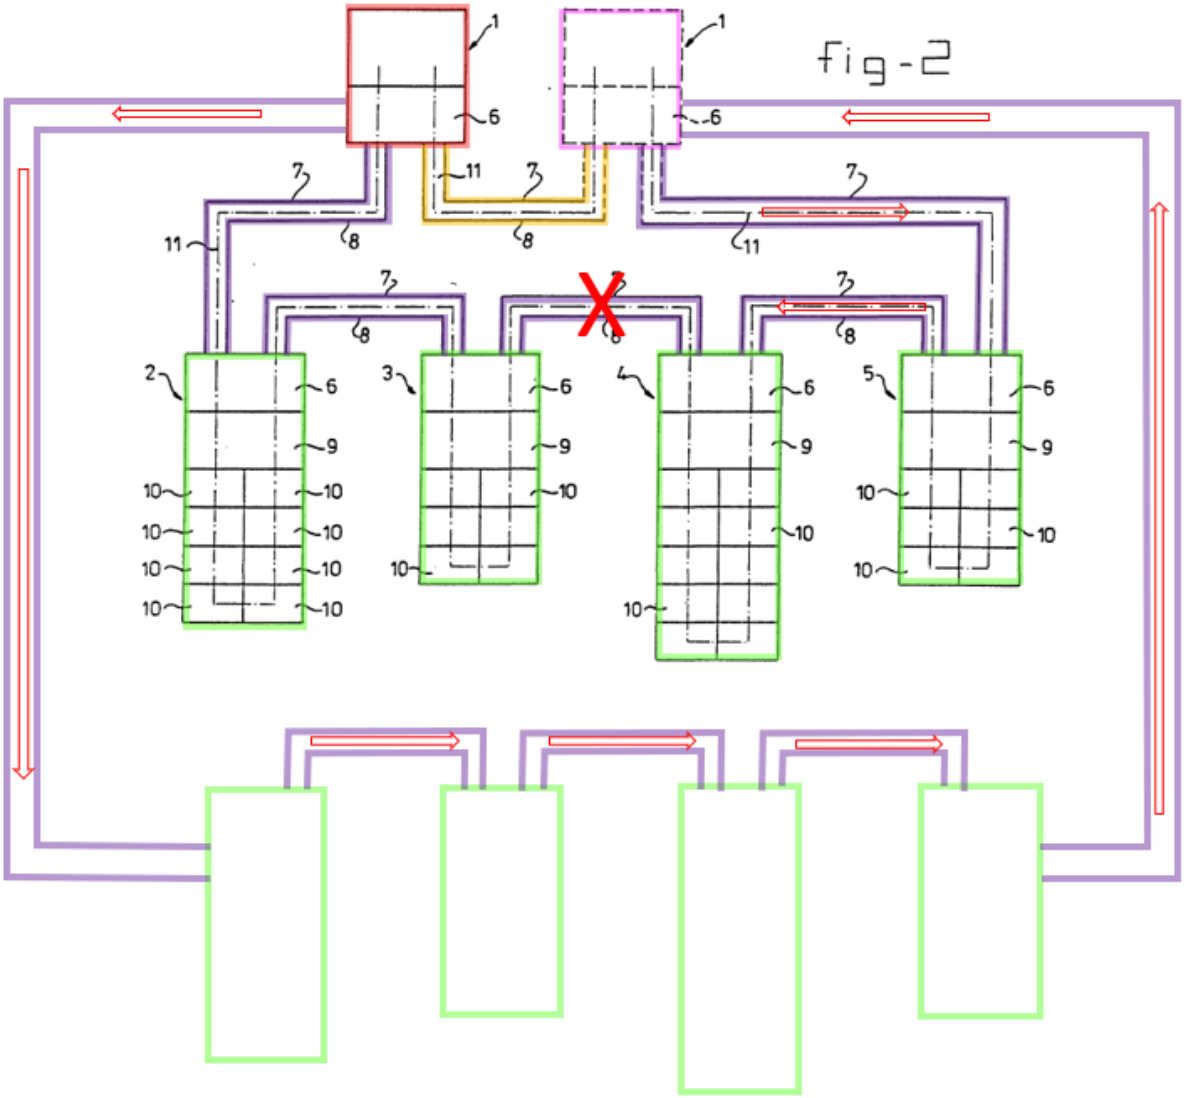
Vink with Patrick (and AAPA for Ground 2) renders obvious claim 4. *See* Section VI.B.1.

b) [7a] wherein in case of a fault at a location in one of the daisy chains, data flow in a portion of the daisy chain between the location of the fault and the second master unit is reversed, so that the downstream data packets are passed from the second master unit to the slave units in the portion of the daisy chain via the last slave unit in the chain, and the upstream data packets are passed by the last slave unit to the second master unit.

As explained in Section VI.B.1.i [9f], Vink with Patrick suggests "wherein in case of a fault at a location in one of the daisy chains, data flow in a portion of the daisy chain between the location of the fault and the second master unit is

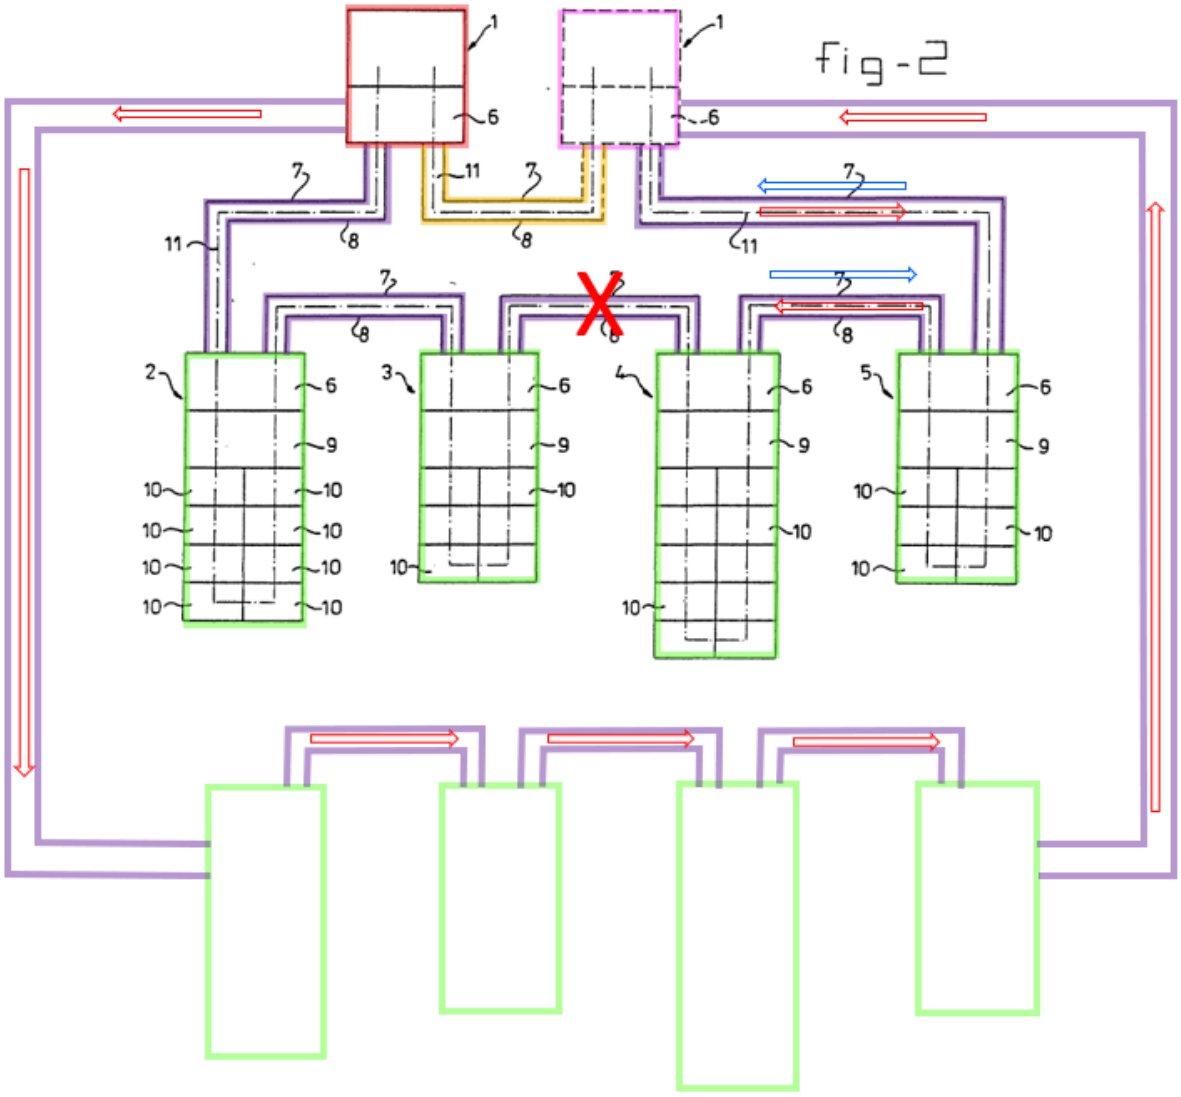
reversed, so that the downstream data packets are passed from the **second master unit** to the slave units in the portion of the daisy chain via the **last slave unit** in the chain."

Specifically, as schematically illustrated below, after re-routing the data to the second daisy chain, a POSITA would have modified Vink's system to allow transmission of downstream data to flow from **slave unit 5** and then to **slave unit 4** in the first daisy chain. Indeed, after having transmitted downstream data across the other daisy chain, transmitting the downstream data to **slave unit 4** through **slave unit 5** is merely the only logical choice to re-routing of data packets, and it would at least have been obvious to try to transmit data in this manner to ensure data reaches as many slave units as possible.



Ex-1005, Figure 2 (modified). And a POSITA would have been able to make such a modification with a reasonable expectation of success because Vink tries "to reach as many connected slaves as possible in a fault situation." Ex-1005, 14:31-34; Ex-1005, 14:29-31 ("[T]he at least one master and slaves are designed to transmit or receive packets of data bits via a transmission system ring in ***one and/or the other direction***."). *See* Ex-1003, ¶173.

Vink also suggests that "the upstream data packets are passed by the **last slave unit** to the **second master unit**." Vink further discloses that the **second master** has the same capabilities as the **first master**. *See* Ex-1005, 16:24-26 ("All the masters are so designed that they can assume the function of the active master, for example in the event of faults."). Accordingly, upon recognizing a fault exists in the daisy chain, a POSITA would have been able to program the **slave units** to send upstream data up to the **second master** (instead of the **first master**) to allow upstream data to be handled by the **second master**, which has the same capabilities as the **first master**, as schematically illustrated below in which the blue arrows indicate the upstream data flow from **slave 4** to the **second master** through the **last slave 5** in the daisy chain.



Ex-1005, Figure 2 (modified). Indeed, Vink teaches that the "passive masters are easily able to detect errors of the active master because … they receive all the information which is intended for the active master." Ex-1005, 16:26-31. This modification merely requires a simple re-routing of data after recognizing faults have occurred and a POSITA would have at least found it obvious to try to make

the proposed modification with a reasonable expectation of success. *See* Ex-1003, ¶174.

5. Dependent Claim 8

a) [8pre] Apparatus according to claim 7,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 7. *See* Section VI.B.4.

b) [8a] wherein the downstream packets for the slave units in the portion of the daisy chain between the location of the fault and the second master unit are conveyed to the second master unit from the first master unit via another one of the daisy chains.

As explained in Sections VI.B.1.i [9f] and VI.B.4.b [7a], Vink with Patrick (and AAPA for Ground 2) suggests [8a].

6. Dependent Claim 10

a) [10pre] Apparatus according to claim 9,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 9. *See* Section VI.B.1.

b) [10a] wherein the pre-switch re-routes the data packets such that substantially no reconfiguration of the switch is required responsive to the fault.

As discussed above in Sections VI.B.1.g [9d], VI.B.1.h [9e], and VI.B.1.i [9f], Vink discloses a **switch** (i.e., **control unit 36** and **transmitting/receiving unit 6**) and a **pre-switch** (i.e., **control computer 38**) in each of the master units.

Vink teaches that when an error occurs, the **control computer 38** "re-routes the data packets." Specifically, although the arrangement of I/O modules within a slave may change, the **control computer 38** would still "provide[] for … the correct sequence of the data bits for a relevant slave, related to the sequence in which the I/O modules are arranged." Ex-1005, 21:4-7. In other words, Vink teaches that the **control computer 38** changes the "calculations" (*see e.g., id.,* 20:35-38) to re-route the data (e.g., provides for the correct sequence of the data bits for a relevant slave, related to the sequence in which the I/O modules are arranged) in view of the error in the transmission system such that "the master is able to reach as many connected slaves as possible in a fault situation." *Id.,* 14:33-34. *See* Ex-1003, ¶179.

This re-routing of the data packets is performed "such that substantially no reconfiguration of the switch [i.e., **control unit 36** and **transmitting/receiving unit 6**] is required." It is the **control computer 38** that "provides for and monitors correct sequence of the data bits for a relevant slave, related to the sequence in which the I/O modules are arranged." Ex-1005, 21:4-7. The **control unit 36** simply "control[s] the data flow via the transmission system 7, 8," (Ex-1005, 20:33-34), based on the newly calculated sequence of the data bits that it "extract[s]" from the control computer 38 and the **transmitting/receiving unit 6** simply "monitor[s]" the transmission. Ex-1005, 21:4-9. Hence, no re-

configuration of the **control unit 36** and **transmitting/receiving unit 6** is required. *See* Ex-1003, ¶180.

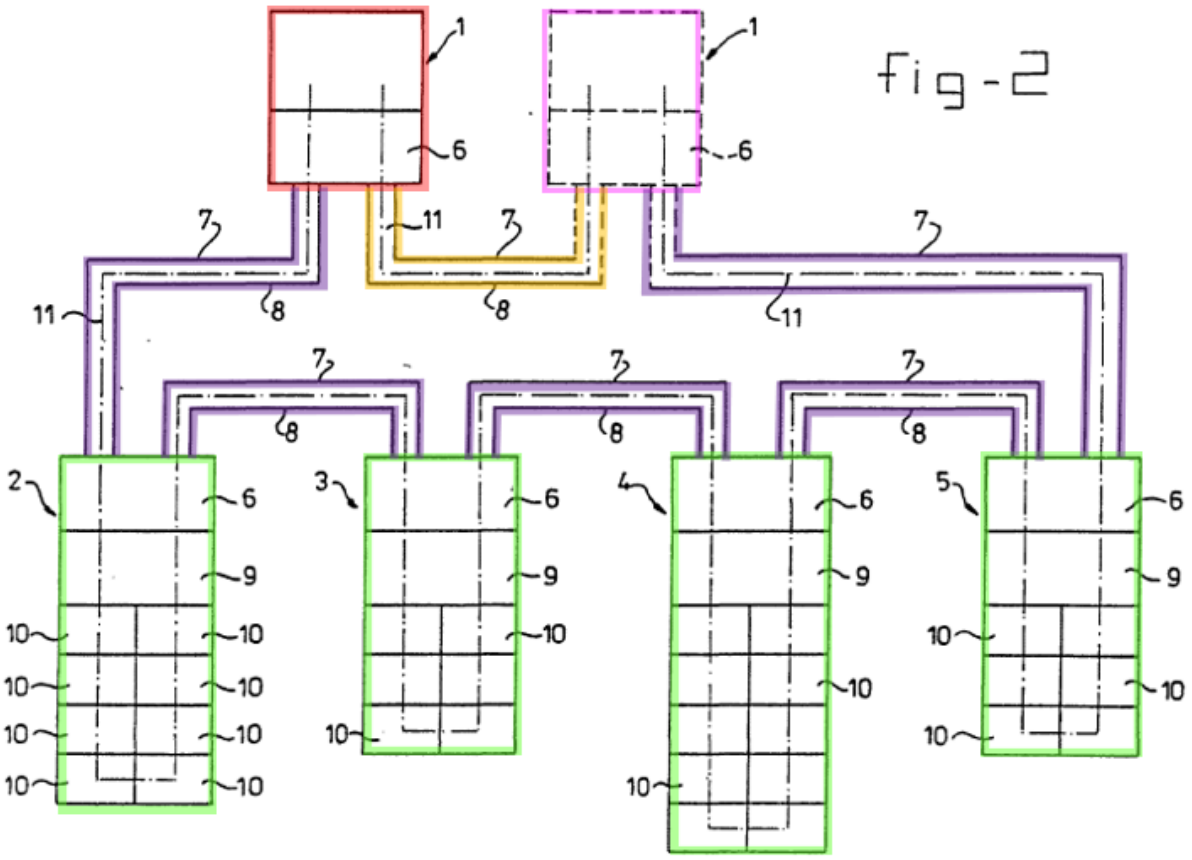
7. Dependent Claim 12

a) [12pre] Apparatus according to claim 11,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 11. *See* Section VI.B.1.

b) [12a] wherein each of the slave units is coupled to receive packets transferred thereto from the first [master unit] and second master unit[] over first and second ones of the physical interface lines, respectively, and

Vink teaches [12a]. For example, Vink teaches that a "very flexible system is obtained … in that the at least one master and slaves are designed to transmit or receive packets of data bits via a transmission system ring in one and/or the other direction." Ex. 1005, 14:28-34. Hence, each of the **slave units 2, 3, 4, and 5** is ***coupled to*** receive packets sent from the **active master 1** and **passive master 1** over the first and second ones of the **transmission lines 7, 8**. *See* Ex-1003, ¶182.



Ex-1005, Figure 2.

c) [12b] wherein the pre-switch passes the packets received through the first and second physical interface line and addressed to any of the ports on the slave unit to respective first and second addresses in the switch fabric.[6]

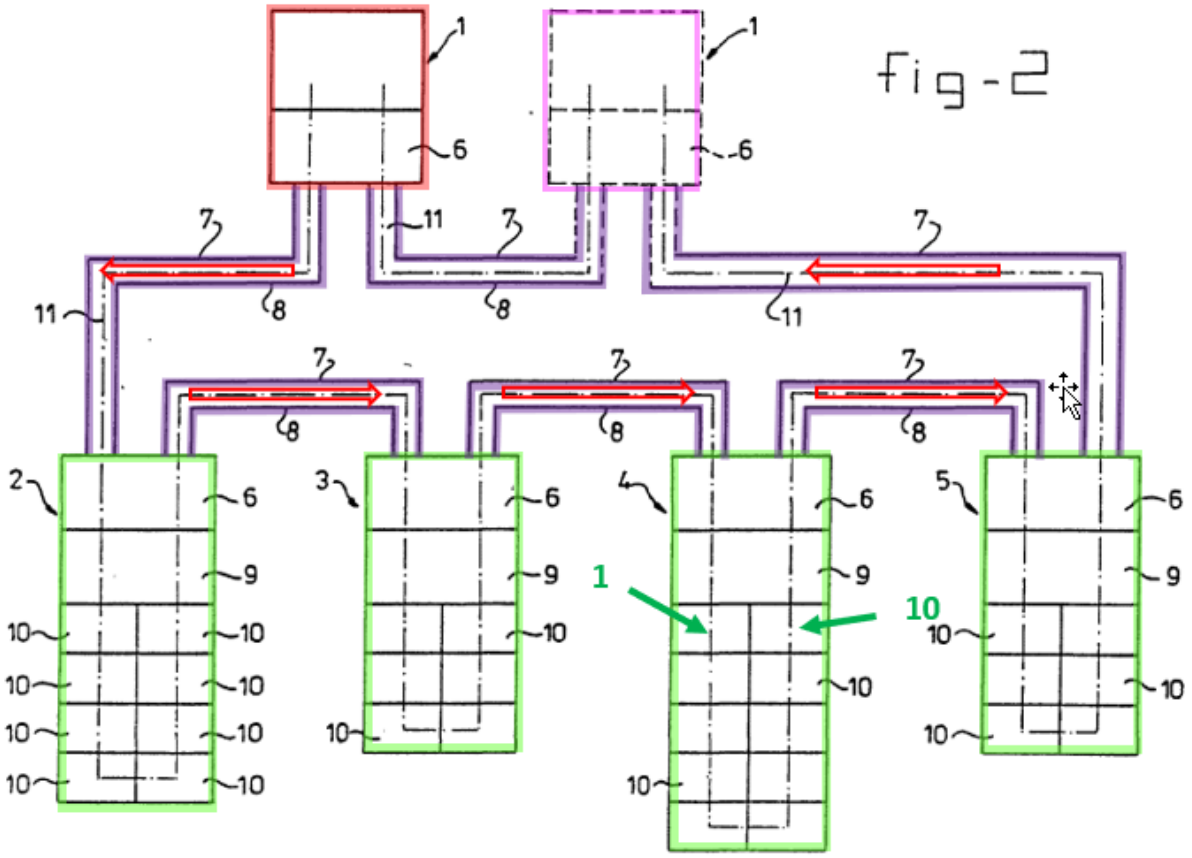
As explained in Sections VI.B.2.b [5a], VI.B.1.k [11e], VI.B.1.l [11f] and VI.B.7.c [12a], Vink discloses [12b]. Specifically, as explained in Sections

[6] Petitioner reserves the right to challenge claim 12 as lacking written description support in other proceedings.

VI.B.1.k [11e] and VI.B.7.b [12a], taking **slave 4** as an exemplary situation, after receiving the downstream data in **slave 4** through *either* of the **transmission lines 7, 8** on the left (under normal operation) or the right (under a fault event) of **slave 4**, the **transmitting/receiving unit 6** and **control module 9** passes the data packets addressed to the respective **ports (e.g., output 33)** on any of the I/O modules 10 within the switch fabric of **slave 4**. *See* Ex-1003, ¶¶183-185.

Vink further discloses that the packets are "addressed to respective first and second addresses in the switch fabric." For example, Vink discloses that "the addressing of the separate I/O modules 10 can take place on the bases of their sequence in the cascade arrangement." Ex-1005, 18:1-3. Vink's addressing is carried out based on "the sequence of the data bits [that] correspond[] to the sequence of the I/O modules." *Id.*, 6:4-5. For example, data packets arriving into the switch fabric of **slave 4** under normal operation would arrive through the **transmission lines 7, 8** on the left side of **slave 4**, in which the I/O modules would be addressed sequentially from, for example, 1 to 10, as shown below. And if a first data packet has as destination the top-left I/O module of **slave 4**, the pre-switch of **slave 4** will receive an address of "1" for the first data packet. If a second data packet has as destination the top-right I/O module of **slave 4**, the pre-switch of **slave 4** will receive an address of "10" for the first data packet, because as illustrated, **slave 4** includes ten I/O modules and because Vink determines the

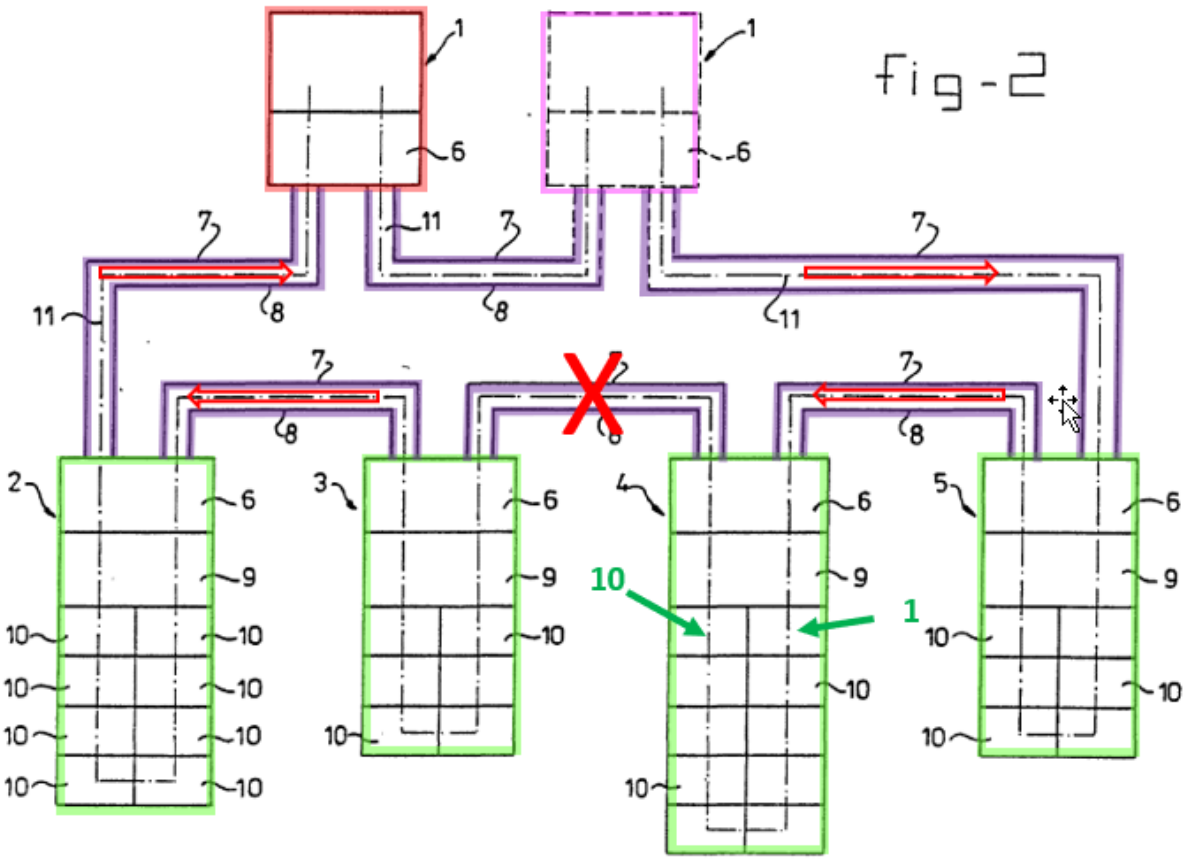
addresses sequentially. Based on the destination addresses of the data packets received at **slave 4**, the pre-switch will pass the data packets to the **slave 4's** switch fabric. *See* Ex-1003, ¶186.



Ex-1005, Figure 2.

In contrast, data packets arriving at the switch fabric of **slave 4** in the event of a fault would arrive through the **transmission lines 7, 8** on the right side of **slave 4**, in which the I/O modules would be addressed sequentially in the opposite order from, for example, 1 to 10, as shown below. *See, e.g.,* Ex-1005, 8:19-20. Specifically, in the event of a fault, the direction of data flow is reversed (as shown

below), the address of the first data packet (with destination for the top-left I/O module of slave 4) that will be received at the pre-switch of slave 4 will be "10," because Vink determines the addresses sequentially within a slave. Similarly, the address of the second data packet (with destination the top-right I/O module of slave 4) that will be received at the pre-switch of slave 4 will be "1." *See* Ex-1003, ¶187.



Ex-1005, Figure 2.

In other words, the destination addresses of the first and second data packets will be swapped at the pre-switch of slave 4 (the address for the first data packet

will be swapped at the pre-switch from 1 to 10 and the address for the second data packet will be swapped at the pre-switch from 10 to 1). Based on the swapped destination addresses the pre-switch will pass the data packets to the **slave 4's** switch fabric. *See* Ex-1003, ¶188.

8. Dependent Claim 13

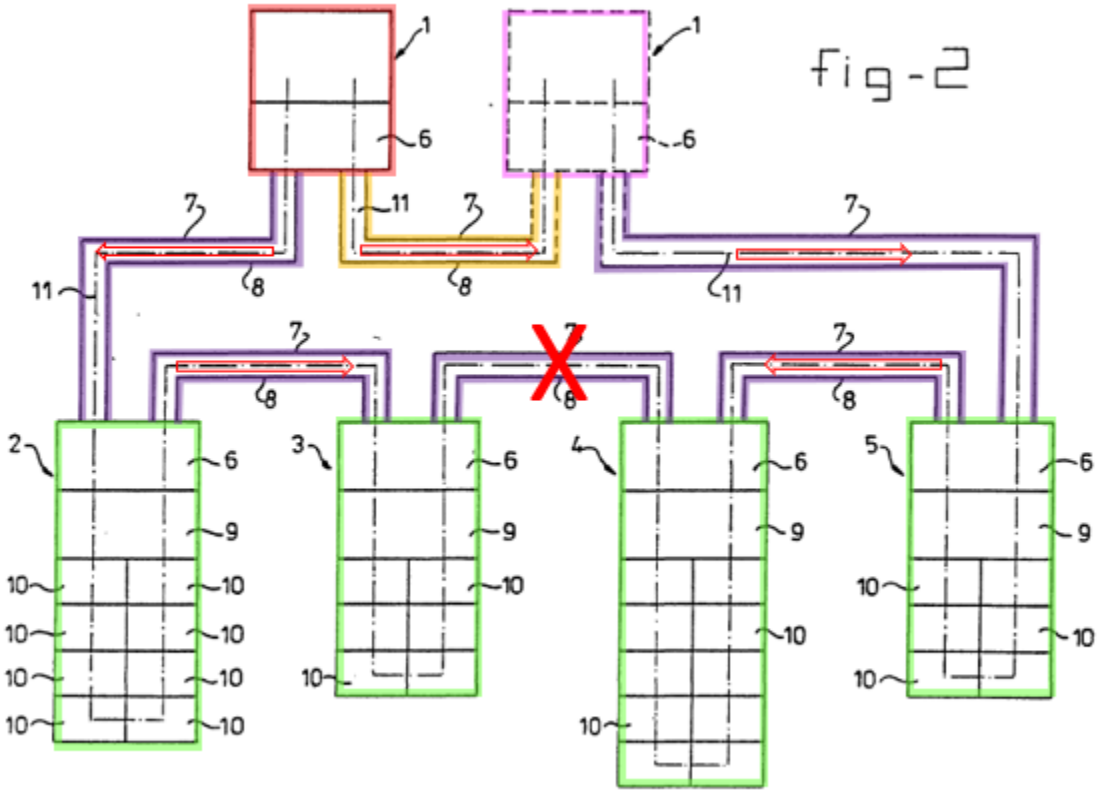
a) [13pre] Apparatus according to claim 12,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 12. *See* Section VI.B.7.

b) [13a] wherein in response to a reversal of a direction of data flow in the daisy chain, the first and second addresses are swapped in the pre-switch, so that substantially no reconfiguration of the switch fabric is required.

Vink discloses [13a]. As explained in Section VI.B.7 [claim 12], Vink teaches that the **control module 9** and **transmitting/receiving unit 6** control that data packets that are delivered to a serially connected I/O modules 10, each of which include **switches** that together form the claimed switch fabric. Ex-1005, 25:1-7 ("the I/O modules (10) being provided with control means (35) and the control module (9) being provided with means (16), coupled to the control means (35), for controlling in a synchronised manner, in response to control information received from the at least one master (1), the inputting and/or outputting of data bits by the connected I/O modules (10)."). *See* Ex-1003, ¶191.

And since the addresses of the I/O modules are based on the sequence of the I/O modules within the sequence of the substations (slaves), reversing the direction of the data flow would thus reverse the sequence of the slaves containing the I/O modules, as well as the sequence of the I/O modules within each slave. And since the destination addresses are swapped at the pre-switch, no reconfiguration at the switch fabric is required at the switch fabric to allow the data packets to reach the appropriate destination. That is, none of the **shift registers 25/26**, **buffers 29/30**, and **control means 35** located in the I/O modules of **slave 4** requires reconfiguration to identify a data packet addressed for a particular I/O module's outputs 33. Thus, when the direction of data flow is switched, Vink discloses swapping the addresses in the pre-switch and no substantial reconfiguration of the switch fabric will be required to route the data packets to the appropriate I/O module. *See* Ex-1003, ¶192.



Ex-1005, Figure 2.

9. **Independent Claim 14**

a) [14pre] In a network access multiplexing system, in which a master unit is connected by a physical interface to a packet-switched network, a slave unit configured to be coupled to the master unit in a daisy chain of such slave units, the slave unit comprising:

To the extent limiting, as explained in Sections VI.B.1.a [1pre/4pre/9pre/11pre], VI.B.1.b [1a/4a/9a/11a], VI.B.1.c [1b/4b/9b/11b], and VI.B.1.d [1c/4c/9c/11c], Vink with Patrick (and AAPA for Ground 2) render obvious [14pre].

And for avoidance of doubt, Vink's communication system is a "multiplexing" system. *See* Ex-1005, 17:5-8 ("The transmitting/receiving unit 6 comprises a multiplexer 12 consisting of two modems 13, 14 which are respectively connected to the first transmission ring 7 and to the second transmission ring 8."), 17:13-17 ("After demodulation by one of the modems 13, 14, the data received are fed to the control module 9 connected to the multiplexer 12 which… is arranged to exchange information with the master in accordance with the known Synchronous Data Link Control (SDLC) protocol."). *See* Ex-1003, ¶¶194-195.

b) [14a] a plurality of ports, for coupling the slave unit to respective subscriber lines;

As explained in Section VI.B.1.c [1b, 4b, 9b, 11b], Vink with Patrick (and AAPA for Ground 2) suggests [14a].

c) [14b] first and second physical interfaces, coupled to exchange packets with preceding and succeeding units, respectively, along the daisy chain;

As explained in Sections VI.B.1.d [1c, 4c, 9c, 11c] and [4d], Vink discloses [14b].

d) [14c] a pre-switch, coupled to receive packets from the first physical interface and responsive to address data carried by the packets, to sort the packets such that packets addressed to the slave unit are retained, and packets

addressed to the succeeding units are passed to the second physical interface; and

As explained in Sections VI.B.1.j [11d] and VI.B.1.l [11f], Vink discloses [14c].

e) [14d] a fabric of one or more switches, which convey the retained packets to the ports, responsive to the address data.

As explained in Sections VI.B.1.j [11d] and VI.B.1.k [11e], Vink discloses [14d].

10. Dependent Claim 15

a) [15pre] The slave unit according to claim 14,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 14. *See* Section VI.B.9.

b) [15a] wherein the pre-switch is further coupled to receive packets transferred thereto from the second physical interface and to sort the packets in like manner to the packets received through the first physical interface.

As explained in Sections VI.B.1.l [11f], VI.B.7.b [12a], VI.B.7.c [12b], and VI.B.8.b [13a], Vink with Patrick suggests [15a].

11. Dependent Claim 16

a) [16pre] The slave unit according to claim 15,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 15. *See* Section VI.B.10.

b) [16a] wherein the retained packets that were received from the first and second physical interfaces and are passed by the pre-switch to the switch fabric are identified by respective first and second port numbers, and

As explained in Section VI.B.7.c [12b], Vink discloses [16a]. To the extent "port numbers" are different from "addresses," a POSITA would have readily recognized that port numbers are often assigned as addresses in a digital communication system. *See e.g.,* U.S. Patent No. 6,631,136 ("Ex-1007"), 13:16-18 ("Destination Address: This field contains the destination address of the Region, Node, ***Port ID***."); *id.,* 20:29-30 ("[A] number of ports contained within each of said nodes, each of said ports being associated with a port address."); U.S. Patent No. 5,425,026 ("Ex-1008"), Abstract ("In a packet switched network, each network node has line and trunk ports to which LAN user terminals and links are connected and which are ***identified by a port address containing a node number plus a port number***"). *See* Ex-1003, ¶203.

c) [16b] wherein in response to a reversal of a direction of data flow in the daisy chain, the first and second port numbers are swapped in the pre-switch, so that substantially no reconfiguration of the switch fabric is required in response to the reversal.

As discussed in Section VI.B.8.b [13a], Vink discloses [16b].

12. Dependent Claim 17

a) [17pre] The slave unit according to claim 14,

Vink with Patrick (and AAPA for Ground 2) suggests claim 14. *See* Section VI.B.9.

b) [17a] wherein when one of the packets received by the pre-switch comprises a multicast packet addressed to one or more of the ports on the slave unit, the pre-switch sorts the multicast packet such that one copy of the packet is retained and another copy of the packet is passed to the second physical interface.

As explained in Sections VI.B.1.j [11d], VI.B.1.l [11f], and VI.B.9.d [14c], Vink discloses [17a].

13. Dependent Claim 18

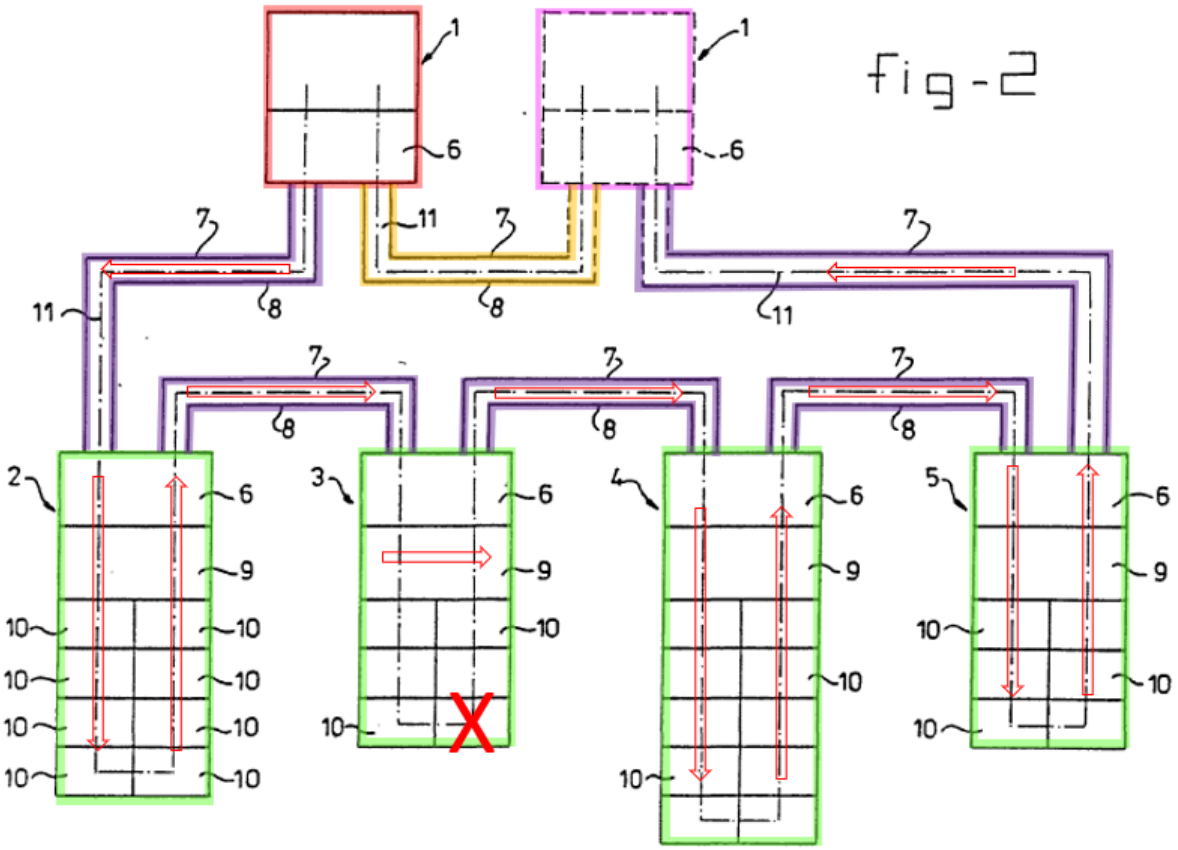
a) [18pre] The slave unit according to claim 14,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 14. *See* Section VI.B.9.

b) [18a] wherein in the event of a fault in the switch fabric, the pre-switch continues to pass the packets addressed to the succeeding units on to the succeeding units without significant interruption.

Vink discloses [18a]. Vink discloses that "[w]hen a transmission error is detected, the control module of a slave will not accept the data received," and the "data bits received are then transferred further unchanged" such that the "master then receives back the data bits d[i]spatched, possibly with transmission errors."

Ex-1005, 14:12-17. In other words, as schematically illustrated below, if one of the I/O modules of slave 3 is determined to be faulty, Vink teaches simply passing the data packets received by the transmitting/receiving unit 6 directly onto the next slave 4 as the control module 9 rejects receipt of that data packets for transmission along the cascade of the I/O modules in slave unit 3. This is possible, for example, but including specific address information bits to the data packets such that data packets can be sent to the specific I/O modules that those data bits were "intended" for without having to through every I/O module. *See also id.*, 8:25-34, 11:29-37, 19:28-32, 21:10-14, 25:20-24. *See* Ex-1003, ¶¶209-210.



14. Independent Claim 19

a) [19pre] A method for providing access to a network, comprising:

Vink "relates to a ***method*** and a communication system for exchanging data in discrete or digital form with one or more input (I) and/or output (O) modules, by means of serial transfer of data bits, under the control of at least one master station (master)." Ex-1005, 1:6-10. Thus, because Vink describes methods for exchanging data with one or more modules, Vink suggests a "*method* for providing

access to a network" as claimed. *See also* Section VI.B.1.a [1pre, 4pre, 9pre, 11pre]; Ex-1003, ¶211.

b) [19a] coupling first and second master units to interface with the network;

As explained in Section VI.B.1.b [1a, 4a, 9a, 11a], Vink with Patrick (and AAPA for Ground 2) suggests [19a].

c) [19b] linking a plurality of slave units, each slave unit comprising one or more ports to respective subscriber lines, in a daisy chain between the first and second master units;

As explained in Sections VI.B.1.c [1b, 4b, 9b, 11b] and VI.B.1.d [1c, 4c, 9c, 11c], Vink with Patrick (and AAPA for Ground 2) suggests [1b], [4b], [9b], and [11b].

d) [19c] conveying initial downstream data packets, received from the network by one of the master units, along the daisy chain in a first direction, so as to deliver the packets to the ports of the slave units; and

As explained in Sections VI.B.1.f [4d], VI.B.1.h [9e], VI.B.1.k [11e], and/or VI.B.1.l [11f], Vink with Patrick (and AAPA for Ground 2) suggests [19c].

e) [19d] in the event of a fault in the daisy chain, conveying further downstream data packets, received from the network by one of the master units, along the daisy chain in a second direction, opposite to the first direction, so

as to deliver the further packets to the ports of at least some of the slave units.

As explained in Sections VI.B.2.b [5a], VI.B.4.b [7a], and/or VI.B.1.i [9f], Vink suggests [19d].

15. Dependent Claim 20

a) [20pre] A method according to claim 19,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 19. *See* Section VI.B.14.

b) [20a] wherein the initial and further downstream packets are received from the network by the first master unit, and

As explained in Sections VI.B.1.f [4d], VI.B.1.h [9e], VI.B.1.k [11e], and/or VI.B.1.l [11f], Vink with Patrick (and AAPA for Ground 2) suggests [20a].

c) [20b] wherein conveying the further downstream packets in the second direction comprises conveying the further downstream packets from the first master unit to the second master unit, and then conveying the further downstream packets from the second master unit to the daisy chain.

As explained in Sections VI.B.2.b [5a], VI.B.4.b [7a], and/or VI.B.1.i [9f], Vink suggests [20b].

16. Dependent Claim 21

a) [21pre] A method according to claim 20,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 20. *See* Section VI.B.15.

b) [21a] wherein conveying the further downstream packets from the first master unit to the second master unit comprises linking further slave units in an additional daisy chain between the first [master unit] and second master unit[], and conveying the further downstream packets from the first master unit to the second master unit over the additional daisy chain.

As explained in Sections VI.B.1.i [9f] and/or VI.B.4.b [7a], Vink with Patrick (and AAPA for Ground 2) suggests [21a].

17. Dependent Claim 22

a) [22pre] A method according to claim 19, and comprising

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 19. *See* Section VI.B.14.

b) [22a] conveying initial upstream data packets, received by the slave units from the subscriber lines, along the daisy chain in the second direction so as to transmit the upstream data packets via the first master unit over the network, and in the event of the fault, conveying further upstream data packets received by one or more of the slave

units along the daisy chain in the first direction via the second master unit.

As explained in Sections VI.B.1.f [4d], VI.B.3.b [6a], and VI.B.4.b [7a], Vink with Patrick suggests [22a].

18. Dependent Claim 23

a) [23pre] A method according to claim 22,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 22. *See* Section VI.B.17.

b) [23a] and comprising bicasting the upstream data packets from the first master unit to the network and to the second master unit, which transmits the bicast upstream data packets over the network.

As explained in Section VI.B.3.b [6a], Vink with Patrick (and AAPA for Ground 2) suggests [23a].

19. Dependent Claim 24

a) [24pre] A method according to claim 19,

Vink with Patrick (and AAPA for Ground 2) renders obvious claim 19. *See* Section VI.B.14.

b) [24a] wherein conveying the initial downstream data packets along the daisy chain comprises pre-switching the packets at each of the slave units, so that packets not addressed to any of the ports on the slave unit are passed to the next slave unit in the daisy chain, while packets that are addressed to one or more of the ports on the slave unit are

passed to a switch fabric that directs the packets to the ports to which they are addressed.

As explained in Section VI.B.1.l [11f], Vink discloses [24a].

20. Dependent Claims 2 and 25

a) [2pre/25pre] Apparatus according to claim 1, / A method according to claim 19,

Vink with Patrick (and AAPA for Ground 2) renders obvious claims 1 and 19. *See* Sections VI.B.1 and VI.B.14.

b) [2a/25a] wherein the network comprises an asynchronous transfer mode (ATM) network.

Vink with Patrick (and AAPA for Ground 2) suggests [2a] and [25a]. Vink discloses a communication system, such as "Local Area Networks" (LAN's). *See* Ex-1005, 1:11-15 ("Communication systems of this type are in practice known per se. Examples are the so-called 'Local Area Networks' (LAN's));" *see also*, *id.*, 1:19-21 ("LAN communication systems are … designed for communicatively coupling two or more connected stations together."). Further, Vink discloses using "public networks" for "long-distance connections." Ex-1005, 1:18-19. *See* Ex-1003, ¶229.

Although Vink does not explicitly teach that the network comprises an ATM network, which can be a public network, and ATM networks were well-known. For example, AAPA discloses that an ATM network could be connected to two

masters. *See* Ex-1001, 1:39-42 ("The master unit comprises a core network interface element 24, providing the necessary physical layer (PHY) and data link layer (for example, ATM) functions."); *see also* Ex-1002, 159 ("[F]igure 2B of the admitted prior art of the instant application is block diagrams that schematically illustrate topologies ***known in the art*** …, ***wherein the core network 22 is an ATM network (packet-switched network)***, page 2, line 1-4." *See* Ex-1003, ¶230.

Patrick similarly discloses that its communication system may be connected to a network, where the network "may be a broadcast network such as ***a local area network ('LAN'), or a non-broadcast circuit-oriented network such as ATM (asynchronous transfer mode)***, frame relay, or x.25." Ex-1006, 5:45-49. *See* Ex-1003, ¶231.

Using a specific protocol like ATM, in lieu of LAN, to allow a communication system to gain access to high-speed network was done routinely in the art and are recognized as known equivalents (as evidenced by Patrick). A POSITA would have had a reasonable expectation of success in making such a modification to Vink's communication system's network connection to provide access to an ATM network. *See* Ex-1003, ¶232.

21. Dependent Claims 3 and 26

a) [3pre/26pre] Apparatus according to claim 1, / A method according to claim 19,

Vink with Patrick (and AAPA for Ground 2) renders obvious claims 1 and 19. *See* Sections VI.B.1 and VI.B.14.

b) [3a/26a] wherein the network comprises an Internet protocol (IP) network.

Vink with Patrick suggests [3a] and [26a]. Vink discloses a communication system, such as "Local Area Networks" (LAN's). *See* Ex-1005, 1:11-15. Although Vink does not explicitly teach that the network comprises an Internet Protocol (IP) network, IP networks were well-known. For example, Patrick discloses that the primary station can be connected to "networks such as the Internet." Ex-1006, 3:20-24. Internet is an Internet Protocol (IP) network. *See* Ex-1003, ¶235.

Using a specific protocol like the Internet Protocol, in lieu of LAN, to allow a communication system to gain access to high-speed network was done routinely in the art and are recognized as known equivalents (as evidenced by Patrick). A POSITA would have had a reasonable expectation of success in making such a modification to Vink's communication system's network connection to provide access to an IP network. *See* Ex-1003, ¶236.

VII. DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE

A. Discretionary Denial Under *Fintiv* Is Not Appropriate

The six factors considered for § 314 denial favor institution. *See Apple Inc. v. Fintiv, Inc*., IPR2020-00019, Paper 11 (Mar. 20, 2020) (precedential).

1. No Evidence Regarding A Stay

No motion to stay has been filed, so this factor is neutral.

2. Parallel Proceeding Trial Date

While trial is currently scheduled for March 4, 2024 (Ex-1010), the projected trial date—based on median time-to-trial statistics—is in August of 2024,"around the same time" as the Board's expected Final Written Decision.[7]

3. Investment In Parallel Proceeding

The co-pending litigation is in its early stages, and the investment in it has been minimal. *See* Ex-1010, Ex-1011. The parties have not exchanged preliminary positions on claim construction, and expert discovery has not begun. *See PEAG LLC v. Varta Microbattery GmbH*, IPR2020-01214, Paper 8, 17 (Jan. 6, 2021). Further, the Markman hearing is not scheduled until September 2023, the expected institution decision date by the Board. Ex-1010, 3-4.

[7] September/October 2024 is 18 months after March/April 2023, when Petitioner expects a notice of accorded filing date for this petition.

4. Overlapping Issues With The Parallel Proceeding

Petitioner stipulates that if the IPR is instituted, Petitioner will not pursue the same grounds in the district court litigation. *Sand Revolution II LLC v. Continental Intermodal Group – Trucking LLC*, IPR2019-01393, Paper 24 at 7 (June 16, 2020) (informative).

5. Identity Of Parties

Petitioner is a defendant in the litigation. This factor should not be a basis for denying institution.

6. Other Circumstances

"[T]he PTAB will not deny institution of an IPR or PGR under *Fintiv* (i) when a petition presents compelling evidence of unpatentability." Memo, 2. Here, the evidence of unpatentability is compelling, and thus the PTAB should not deny institution under *Fintiv*.

B. Discretionary Denial Under 35 U.S.C. § 325(d) Is Not Appropriate

This petition challenges each of the claims using a combination of Vink, Patrick and/or AAPA. While AAPA was applied during the prosecution of the '904 Patent, the limitations that were deemed allowable are taught by Vink and/or Patrick. The Examiner erred by overlooking the teachings of Vink and Patrick, neither of which were considered by the Examiner. Ex-1001, 1. Discretionary denial is therefore not appropriate. *Advanced Bionics, LLC v. MED-EL*

Elektromedizinische Geräte GmbH, IPR2019-01469, Paper 6 (Feb. 13, 2020) (precedential). Accordingly, the Board should not exercise its discretion to deny this petition under § 325(d).

C. Discretionary Denial Under *General Plastic* Is Not Appropriate

The '904 patent has not been challenged in any prior IPR petition, so none of *General Plastic* discretionary institution factors apply to this Petition.

VIII. CONCLUSION

Accordingly, Petitioner has established a reasonable likelihood that the Challenged Claims are unpatentable.

Respectfully submitted,

Dated: March 14, 2023
Desmarais LLP
230 Park Ave
New York, NY 10023

/Yung-Hoon Ha/
Yung-Hoon Ha
Lead Counsel for Petitioner
Registration No. 56,368

IX. 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 '904 patent is or was involved in the following case:

Case Heading	Number	Court	Date
Orckit Corporation v. Cisco Systems, Inc.	2-22-cv-00276	EDTX	Jul. 7, 2022

C. Lead and Back-up Counsel and Service Information

Lead Counsel
Yung-Hoon Ha
Desmarais LLP
230 Park Ave
New York, NY 10023

Phone: (212) 351-3411
yha@desmaraisllp.com
USPTO Reg. No. 56,368

Back-up Counsels
Theodoros Konstantakopoulos
Desmarais LLP
230 Park Ave
New York, NY 10023

Phone: (212) 351-3411
tkonstantakopoulos@desmaraisllp.com
USPTO Reg. No. 74,155

Emily Weber
Desmarais LLP
101 California Street
San Francisco, CA 94111

Phone: (415) 573-1858
eweber@desmaraisllp.com
USPTO Reg. No. 79,973

Please address all correspondence to lead and back-up counsel. Petitioner consents to service in this proceeding by email at CiscoOrckitIPRService@desmaraisllp.com and the email addresses above.

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioner hereby certifies, in accordance with and in 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 13,994. Pursuant to 37 C.F.R. § 42.24(d), this word count excludes the table of contents, mandatory notices under § 42.8, certificate of service, certificate of word count, and appendix of exhibits.

Dated: March 14, 2023

/Yung-Hoon Ha/
Yung-Hoon Ha
Lead Counsel for Petitioner
Registration No. 56,368

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 at the correspondence address of record, as detailed below.

Date of service March 14, 2023

Manner of service Federal Express Mail

Documents served
- Petition for *Inter Partes* Review of U.S. 6,680,904
- Petitioner's Exhibit List
- Exhibits 1001-1011
- Certificate of Word Count
- Petitioner's Power of Attorney

Persons served
May Patent Ltd.
c/o Dorit Shem-Tov
P.O. Box 7230
Ramat-Gan, 5217102
Israel

/Yung-Hoon Ha/
Yung-Hoon Ha
Lead Counsel for Petitioner
Registration No. 56,368

EXHIBIT 5

Exhibit A: Asserted Claims by Product Category (as amended on 1/19/2023)

The following is a list of claims asserted by the Cisco accused products that infringe on U.S. Patents No. 7,545,740, 8,830,821, 6,680,904 and 10,652,111.

Claims	Product Categories
'740 Patent – Claims 1-31	• Cisco ASR 900 Series ○ Cisco ASR 901 Router ○ Cisco ASR 901S ○ Cisco ASR 902 Router ○ Cisco ASR 902U Router ○ Cisco ASR 903 Router ○ Cisco ASR 903U Router ○ Cisco ASR 907 Router ○ Cisco ASR 914 Router • Cisco IOS XRv 9000 Router • Cisco CRS-1 8-Slot Single Shelf System • Cisco CRS-1 16-Slot Single Shelf System • Cisco CRS-3 8-Slot Single Shelf System • Cisco CRS-3 16-Slot Single Shelf System • Cisco CRS-X 8-Slot Single-Shelf System • Cisco CRS-X 16-Slot Single-Shelf System • Cisco CRS-X Multishelf System • Cisco 12000 Series Routers ○ 12004 Router ○ 12006 Router ○ 12010 Router ○ 12016 Router ○ 12404 Router ○ 12406 Router ○ 12410 Router ○ 12416 Router ○ 12810 Router ○ 12816 Router • Cisco 8000 Series Routers ○ Cisco 8100 Series Routers ▪ 8101 ▪ 8102 ▪ 8111 ○ Cisco 8200 Series Routers

Claims	Product Categories
	▪ 8201 ▪ 8202 ○ Cisco 8800 Series Routers ▪ 8804 ▪ 8808 ▪ 8812 ▪ 8818 • Cisco 7600 Series Router ○ 7603 ○ 7604 ○ 7606 ○ 7609 ○ 7613 • Cisco ASR 920 Series Routers (all models) • Cisco ASR 1000 Router ○ ASR 1001 ○ ASR 1002 ○ ASR 1004 ○ ASR 1006 ○ ASR 1009 ○ ASR 1013 • Cisco ASR 9000 Series ○ ASR 9001 Router ○ ASR 9006 Router ○ ASR 9010 Router ○ ASR 9901 Router ○ ASR 9902 Router ○ ASR 9903 Router ○ ASR 9904 Router ○ ASR 9906 Router ○ ASR 9910 Router ○ ASR 9912 Router ○ ASR 9922 Router ○ ASR 9000v-V2 • Cisco 4000 Series ISR ○ ISR 4221 Router ○ ISR 4331 Router ○ ISR 4431 Router ○ ISR 4461 Router • Cisco 800 Series ○ 800 Series Industrial ISR Router ○ 800M Series ○ 810 Router

Claims	Product Categories
	○ 860 Router ○ 880 Router ○ 890 Router • Cisco ISR 900 ○ ISR 921 ○ ISR 926 ○ ISR 927 ○ ISR 931 • Cisco 1000 ISR ○ ISR 1100 ○ ISR 1101 ○ ISR 1109 ○ ISR 111x ○ ISR 1111X ○ ISR 1120 ○ ISR 1131 ○ ISR 1160 • Cisco Cloud Services Router 1000V Series • Cisco Catalyst IR1100 Rugged Series Routers • Cisco Catalyst 1000 Series ○ Compact form-factor models ○ General-purpose models ○ Fast Ethernet models • Cisco Catalyst 6500 Series ○ Catalyst 6503 ○ Catalyst 6504 ○ Catalyst 6506 ○ Catalyst 6509 ○ Catalyst 6513 • Cisco Catalyst 8200 Series ○ C8200-1N-4T ○ C8200L-1N-4T • Cisco Catalyst 8300 Series ○ C8300-1N1S-6T ○ C8300-1N1S-4T2X ○ C8300-2N2S-6T ○ C8300-2N2S-4T2X • Cisco Catalyst 8500 Series ○ C8500-12X4QC ○ C8500-12X ○ C8500L-8S4X

Claims	Product Categories
	• Cisco Catalyst 9200 Series ○ Catalyst 9200 enhanced VN ○ Catalyst 9200 multigigabit ○ Catalyst 9200 1G ○ Catalyst 9200L multigigabit ○ Catalyst 9200L 1G ○ Catalyst 9200CX compact • Cisco Catalyst 9300 Series ○ Catalyst 9300X copper ○ Catalyst 9300X fiber ○ Catalyst 9300 high-performance ○ Catalyst 9300 UPOE+ ○ Catalyst 9300 1G ○ Catalyst 9300L/LM 1G • Cisco Catalyst 9400 Series ○ Catalyst 9400X SUP-2/2XL ○ Catalyst 9400 SUP-1/1XL/1XL-Y ○ Catalyst 9400X line cards ○ Multigigabit/UPOE line cards ○ Copper line cards • Cisco Catalyst 9500 Series ○ Catalyst 9500X 400G, 100G/40G switches ○ Catalyst 9500 100G high-performance switches ○ Catalyst 9500 40G high-performance switches ○ Catalyst 9500 25G high-performance switches ○ Catalyst 9500 40G switches ○ Catalyst 9500 10G/1G switches • Cisco Catalyst 9600 Series ○ Catalyst 9600X Supervisor Engine 2 ○ Catalyst 9600X 100G/40G 400G/200G line card ○ Catalyst 9600 Supervisor Engine 1 ○ Catalyst 9600 50G 400G/200G line card ○ Catalyst 9600 100G/40G line card ○ Catalyst 9600 25G/10G/1G line card ○ Catalyst 9600 1G line card

Claims	Product Categories
	○ Catalyst 9600 Series multigigabit line card • Cisco IE 2000U Series Switches ○ Catalyst 9600 Series multigigabit line card ○ IE-2000U-4TS-G ○ IE-2000U-4T-G ○ IE-2000U-4S-G ○ IE-2000U-8TC-G ○ IE-2000U-16TC-G ○ IE-2000U-16TC-G-X ○ IE-2000U-16TC-GP • Cisco ME 4900 Series Switches • Cisco 2500 Series Connected Grid Switches ○ CGS-2520-16S-8PC Connected Grid Switch ○ CGS-2520-24TC Connected Grid Switch • Cisco Catalyst IR1100 • Cisco Catalyst IR1800 ○ IR1821-K9 ○ IR1831-K9 ○ IR1833-K9 ○ IR1835-K9 • Cisco Network Convergence System 4000 ○ NCS 4009 ○ NCS 4016 • Cisco Network Convergence System 4200 ○ NCS 4201 ○ NCS 4202 ○ NCS 4206 ○ NCS 4216 ○ NCS 4216 (F2B System) • Cisco Network Convergence System 5700 ○ Line Cards: ▪ NC57-18DD-SE ▪ NC57-24DD ▪ NC57-36H-SE ▪ NC57-36H6D-S

Claims	Product Categories
	○ Routers: ▪ NCS 57B1-6D24-SYS /NCS 57B1-5DSE-SYS ▪ NCS 57C3-MOD-SYS /NCS 57C3-MODS-SYS ▪ NCS 57C1-48Q6-SYS • Cisco Network Convergence System 5000 ○ ENCS 5100 model ○ ENCS 5400 model ○ NCS 5001 ○ NCS 5002 • Cisco Network Convergence System 5500 ○ NCS 5501 ○ NCS 5501-SE ○ NCS 5502 ○ NCS 5502-SE ○ NCS 5508 ○ NCS 5516 • Cisco Network Convergence System 6000 • Cisco Network Convergence System 500 ○ Network Convergence System 520 Routers ○ Network Convergence System 540 Routers ○ Network Convergence System 560 Routers • Cisco Carrier Packet Transport 600 • Cisco Carrier Packet Transport 200 • Cisco Nexus 3000 Series Switches ○ Nexus 3016 Switch ○ Nexus 3048 Switch ○ Nexus 3064 Switch ○ Nexus 3064-T Switch ○ Nexus 3132C-Z Switch ○ Nexus 3132Q Switch ○ Nexus 3132Q-V Switch ○ Nexus 3132Q-XL Switch ○ Nexus 3164Q Switch ○ Nexus 3172PQ Switch

Claims	Product Categories
	○ Nexus 3172PQ-XL Switch ○ Nexus 3172TQ Switch ○ Nexus 3172TQ-32T Switch ○ Nexus 3172TQ-XL Switch ○ Nexus 3232C Switch ○ Nexus 3264C-E Switch ○ Nexus 3264Q Switch ○ Nexus 3408-S Switch ○ Nexus 3432D-S Switch ○ Nexus 3464C Switch ○ Nexus 3524 Switch ○ Nexus 3524-X Switch ○ Nexus 3524-XL Switch ○ Nexus 3548 Switch ○ Nexus 3548-X Switch ○ Nexus 3548-XL Switch ○ Nexus 31108PC-V Switch ○ Nexus 31108TC-V Switch ○ Nexus 31128PQ Switch ○ Nexus 34180YC Switch ○ Nexus 34200YC-SM Switch ○ Nexus 36180YC-R Switch • Cisco Nexus 7000 Series Switches ○ Nexus 7000 ○ Nexus 7700 • Cisco Nexus 9000 Series Switches ○ Nexus 9000v Switch ○ Nexus 9236C Switch ○ Nexus 9272Q Switch ○ Nexus 9316D-GX Switch ○ Nexus 9332C Switch ○ Nexus 9332D-GX2B Switch ○ Nexus 9332PQ Switch ○ Nexus 9336C-FX2 Switch ○ Nexus 9336C-FX2-E Switch ○ Nexus 9336PQ ACI Spine Switch ○ Nexus 9348D-GX2A Switch ○ Nexus 9348GC-FXP Switch ○ Nexus 9364C Switch ○ Nexus 9364C-GX Switch ○ Nexus 9364D-GX2A Switch ○ Nexus 9372PX Switch ○ Nexus 9372PX-E Switch ○ Nexus 9372TX Switch

Claims	Product Categories
	○ Nexus 9372TX-E Switch ○ Nexus 9396PX Switch ○ Nexus 9396TX Switch ○ Nexus 9508 Switch ○ Nexus 9808 Switch ○ Nexus 92160YC-X Switch ○ Nexus 92300YC Switch ○ Nexus 92304QC Switch ○ Nexus 92348GC-X Switch ○ Nexus 93108TC-EX Switch ○ Nexus 93108TC-EX-24 Switch ○ Nexus 93108TC-FX3H Switch ○ Nexus 93108TC-FX3P Switch ○ Nexus 93108TC-FX Switch ○ Nexus 93108TC-FX-24 Switch ○ Nexus 93120TX Switch ○ Nexus 93128TX Switch ○ Nexus 93180LC-EX Switch ○ Nexus 93180YC-EX Switch ○ Nexus 93180YC-EX-24 Switch ○ Nexus 93180YC-FX3 Switch ○ Nexus 93180YC-FX3H Switch ○ Nexus 93180YC-FX3S Switch ○ Nexus 93180YC-FX Switch ○ Nexus 93180YC-FX-24 Switch ○ Nexus 93216TC-FX2 Switch ○ Nexus 93240YC-FX2 Switch ○ Nexus 93360YC-FX2 Switch ○ Nexus 93600CD-GX Switch • Meraki MX64 ○ MX64W Switch • Meraki MX67 ○ MX67C Switch ○ MX67W Switch • Meraki MX68 ○ MX68C Switch ○ MX68W Switch • Meraki MX75
’821 Patent – Claims 1-20	• Cisco 7200 Series Routers ○ 7201 ○ 7202 ○ 7204 ○ 7204VXR ○ 7206

Claims	Product Categories
	○ 7206VXR • Cisco 12000 Series Routers ○ 12004 Router ○ 12006 Router ○ 12010 Router ○ 12016 Router ○ 12404 Router ○ 12406 Router ○ 12410 Router ○ 12416 Router ○ 12810 Router ○ 12816 Router • Cisco ASR 900 Series Routers ○ ASR 901 Router ○ ASR 901S ○ ASR 902 Router ○ ASR 902U Router ○ ASR 903 Router ○ ASR 903U Router ○ ASR 907 Router ○ ASR 914 Router • Cisco ASR 920 Series Routers (all models) • Cisco ASR 1000 Series Routers ○ ASR 1001 ○ ASR 1002 ○ ASR 1004 ○ ASR 1006 ○ ASR 1009 ○ ASR 1013 • Cisco ASR 9000 Series Aggregation Services Routers ○ ASR 9001 Router ○ ASR 9006 Router ○ ASR 9010 Router ○ ASR 9901 Router ○ ASR 9902 Router ○ ASR 9903 Router ○ ASR 9904 Router ○ ASR 9906 Router ○ ASR 9910 Router ○ ASR 9912 Router ○ ASR 9922 Router ○ ASR 9000v-V2

Claims	Product Categories
	• Cisco Catalyst 3750 Series Switches • Catalyst 3750V2-24FS Switch • Catalyst 3750V2-24PS Switch • Catalyst 3750V2-24TS Switch • Catalyst 3750V2-48PS Switch • Catalyst 3750V2-48TS Switch • Catalyst 3750G-12S Switch • Catalyst 3750G-12S-SD Switch • Catalyst 3750G-24PS Switch • Catalyst 3750G-24T Switch • Catalyst 3750G-24TS Switch • Catalyst 3750G-24TS-1U Switch • Catalyst 3750G-48TS Switch • Catalyst 3750G-48PS Switch • Catalyst 3750-24FS Switch • Catalyst 3750G-24WS Switch • Catalyst 3750-24TS Switch • Catalyst 3750-48TS Switch • Catalyst 3750-24PS Switch • Catalyst 3750-48PS Switch • Catalyst 3750G-16TD Switch • Cisco Catalyst 6500 Series Switches ○ Catalyst 6503 ○ Catalyst 6504 ○ Catalyst 6506 ○ Catalyst 6509 ○ Catalyst 6513 • Cisco Catalyst 6800 Series Switches ○ Catalyst 6807-XL Switch ○ Catalyst 6840-X Switch ○ Catalyst 6880-X Switch ○ Catalyst C6816-X-LE Switch ○ Catalyst C6824-X-LE-40G Switch ○ Catalyst C6832-X-LE Switch ○ Catalyst C6840-X-LE-40G Switch ○ Catalyst 6800ia • Cisco Catalyst 8200 Series Edge Platforms ○ C8200-1N-4T ○ C8200L-1N-4T • Cisco Catalyst 8200 Series Edge Ucpe ○ C8200-UCPE-1N8

Claims	Product Categories
	• Cisco Catalyst 8300 Series Edge Platforms ○ C8300-1N1S-4T2X ○ C8300-1N1S-6T ○ C8300-2N2S-4T2X ○ C8300-2N2S-6T • Cisco Catalyst 8500 Series Switches ○ C8500-12X4QC ○ C8500-12X ○ C8500L-8S4X • Cisco Catalyst 9200 Series Switches ○ Catalyst 9200 enhanced VN ○ Catalyst 9200 multigigabit ○ Catalyst 9200 1G ○ Catalyst 9200L multigigabit ○ Catalyst 9200L 1G ○ Catalyst 9200CX compact • Cisco Catalyst 9300 Series Switches ○ Catalyst 9300 ○ Catalyst 9300 High Performance ○ Catalyst 9300L ○ Catalyst 9300LM ○ Catalyst 9300X ○ Catalyst 9300 UPOE+ ○ Catalyst 9300 1G • Cisco Catalyst 9400 Series Switches ○ Catalyst 9400X SUP-2/2XL ○ Catalyst 9400 SUP-1/1XL/1XL-Y ○ Catalyst 9400X line cards ○ Multigigabit/UPOE line cards ○ Copper line cards • Cisco Catalyst 9500 Series Switches ○ Catalyst 9500X 400G, 100G/40G switches ○ Catalyst 9500 100G high-performance switches ○ Catalyst 9500 40G high-performance switches ○ Catalyst 9500 25G high-performance switches ○ Catalyst 9500 40G switches ○ Catalyst 9500 10G/1G switches • Cisco Catalyst 9600 Series Switches ○ Catalyst 9600X Supervisor Engine 2

Claims	Product Categories
	○ Catalyst 9600X 100G/40G 400G/200G line card ○ Catalyst 9600 Supervisor Engine 1 ○ Catalyst 9600 50G 400G/200G line card ○ Catalyst 9600 100G/40G line card ○ Catalyst 9600 25G/10G/1G line card ○ Catalyst 9600 1G line card ○ Catalyst 9600 Series multigigabit line card • Cisco Catalyst ESS9300 Embedded Series Switches • Cisco Catalyst IR1101 Rugged Series Routers • Cisco Catalyst IR1800 Rugged Series Routers ○ IR1821-K9 ○ IR1831-K9 ○ IR1833-K9 ○ IR1835-K9 • Cisco Catalyst IR 8300 Rugged Series • Cisco Cloud Services Router 1000V Series • Cisco 1000 Series Connected Grid Routers ○ 1240 Connected Grid Router ○ 1120 Connected Grid Router • Cisco 2000 Series Connected Grid Routers • Cisco 5900 Series Embedded Services Routers ○ 5921 Embedded Services Router ○ 5940 Embedded Services Router • Cisco 800 Series Industrial Integrated Services Routers ○ 829 Integrated Services Routers ○ 807 Integrated Services Routers ○ 809 Integrated Services Routers • Cisco 900 Series Integrated Services Routers ○ ISR 921 ○ ISR 926 ○ ISR 927

Claims	Product Categories
	○ ISR 931 • Cisco 1000 Series Integrated Services Routers ○ ISR 1100 ○ ISR 1101 ○ ISR 1109 ○ ISR 111x ○ ISR 1111X ○ ISR 1120 ○ ISR 1131 ○ ISR 1160 • Cisco 4000 Integrated Services Routers • Cisco Network Convergence Systems 5000 Series • Cisco Network Convergence Systems 4000 Series • Cisco Network Convergence Systems 540 Series Routers ○ ENCS 5100 model ○ ENCS 5400 model ○ NCS 5001 ○ NCS 5002 • Cisco Network Convergence Systems 4200 Series ○ NCS 4201 ○ NCS 4202 ○ NCS 4206 ○ NCS 4216 ○ NCS 4216 (F2B System) • Cisco Network Convergence System 5500 Series ○ NCS 5501 ○ NCS 5501-SE ○ NCS 5502 ○ NCS 5502-SE ○ NCS 5508 ○ NCS 5516 • Cisco Network Convergence Systems 500 Series ○ Network Convergence System 520 Routers ○ Network Convergence System 540 Routers

Claims	Product Categories
	○ Network Convergence System 560 Routers • Cisco Network Convergence Systems 5700 Series ○ Line Cards: ▪ NC57-18DD-SE ▪ NC57-24DD ▪ NC57-36H-SE ▪ NC57-36H6D-S ○ Routers: ▪ NCS 57B1-6D24-SYS /NCS 57B1-5DSE-SYS ▪ NCS 57C3-MOD-SYS /NCS 57C3-MODS-SYS ▪ NCS 57C1-48Q6-SYS • Cisco Cloud Native Broadband Router • Cisco Carrier Routing Systems • Cisco Network Convergence Systems 6000 Series Routers • Cisco 8000 Series Routers ○ Cisco 8100 Series Routers ▪ 8101 ▪ 8102 ▪ 8111 ○ Cisco 8200 Series Routers ▪ 8201 ▪ 8202 ○ Cisco 8800 Series Routers ▪ 8804 ▪ 8808 ▪ 8812 ▪ 8818 • Cisco Nexus 3000 Series Switches ○ Nexus 3016 Switch ○ Nexus 3048 Switch ○ Nexus 3064 Switch ○ Nexus 3064-T Switch ○ Nexus 3132C-Z Switch ○ Nexus 3132Q Switch ○ Nexus 3132Q-V Switch ○ Nexus 3132Q-XL Switch ○ Nexus 3164Q Switch ○ Nexus 3172PQ Switch ○ Nexus 3172PQ-XL Switch

Claims	Product Categories
	○ Nexus 3172TQ Switch ○ Nexus 3172TQ-32T Switch ○ Nexus 3172TQ-XL Switch ○ Nexus 3232C Switch ○ Nexus 3264C-E Switch ○ Nexus 3264Q Switch ○ Nexus 3408-S Switch ○ Nexus 3432D-S Switch ○ Nexus 3464C Switch ○ Nexus 3524 Switch ○ Nexus 3524-X Switch ○ Nexus 3524-XL Switch ○ Nexus 3548 Switch ○ Nexus 3548-X Switch ○ Nexus 3548-XL Switch ○ Nexus 31108PC-V Switch ○ Nexus 31108TC-V Switch ○ Nexus 31128PQ Switch ○ Nexus 34180YC Switch ○ Nexus 34200YC-SM Switch ○ Nexus 36180YC-R Switch • Cisco Nexus 7000 Series Switches ○ Nexus 7000 ○ Nexus 7700 • Cisco Nexus 9000 Series Switches ○ Nexus 9000v Switch ○ Nexus 9236C Switch ○ Nexus 9272Q Switch ○ Nexus 9316D-GX Switch ○ Nexus 9332C Switch ○ Nexus 9332D-GX2B Switch ○ Nexus 9332PQ Switch ○ Nexus 9336C-FX2 Switch ○ Nexus 9336C-FX2-E Switch ○ Nexus 9336PQ ACI Spine Switch ○ Nexus 9348D-GX2A Switch ○ Nexus 9348GC-FXP Switch ○ Nexus 9364C Switch ○ Nexus 9364C-GX Switch ○ Nexus 9364D-GX2A Switch ○ Nexus 9372PX Switch ○ Nexus 9372PX-E Switch ○ Nexus 9372TX Switch ○ Nexus 9372TX-E Switch

Claims	Product Categories
	○ Nexus 9396PX Switch ○ Nexus 9396TX Switch ○ Nexus 9508 Switch ○ Nexus 9808 Switch ○ Nexus 92160YC-X Switch ○ Nexus 92300YC Switch ○ Nexus 92304QC Switch ○ Nexus 92348GC-X Switch ○ Nexus 93108TC-EX Switch ○ Nexus 93108TC-EX-24 Switch ○ Nexus 93108TC-FX3H Switch ○ Nexus 93108TC-FX3P Switch ○ Nexus 93108TC-FX Switch ○ Nexus 93108TC-FX-24 Switch ○ Nexus 93120TX Switch ○ Nexus 93128TX Switch ○ Nexus 93180LC-EX Switch ○ Nexus 93180YC-EX Switch ○ Nexus 93180YC-EX-24 Switch ○ Nexus 93180YC-FX3 Switch ○ Nexus 93180YC-FX3H Switch ○ Nexus 93180YC-FX3S Switch ○ Nexus 93180YC-FX Switch ○ Nexus 93180YC-FX-24 Switch ○ Nexus 93216TC-FX2 Switch ○ Nexus 93240YC-FX2 Switch ○ Nexus 93360YC-FX2 Switch ○ Nexus 93600CD-GX Switch
’111 Patent – Claims 1-9, 12-24, and 27-31	• Cisco SD-WAN • Cisco Catalyst 8500 Series Edge Platforms ○ C8500-12X ○ C8500-12X4QC ○ C8500L-8S4X • Cisco Catalyst 8300 Series Edge Platforms ○ C8300-1N1S-4T2X ○ C8300-1N1S-6T ○ C8300-2N2S-4T2X ○ C8300-2N2S-6T • Cisco Catalyst 8200 Series Edge Platforms ○ C8200-1N-4T

Claims	Product Categories
	○ C8200L-1N-4T • Cisco Catalyst 8200 uCPE Series Edge Platforms ○ C8200-UCPE-1N8 • Cisco ASR 1000 Series Aggregation Services Routers ○ ASR 1001-HX ○ ASR 1001-X ○ ASR 1002-HX ○ ASR 1002-X ○ ASR 1006-X • Cisco ISR 4000 Series Integrated Services Routers ○ ISR 4321 ○ ISR 4331 ○ ISR 4351 ○ ISR 4221 ○ ISR 4221X ○ ISR 4431 ○ ISR 4451 ○ ISR 4461 • Cisco ISR 1100 Series Integrated Services Routers ○ ISR1100-4G ○ ISR1100-4GLTE ○ ISR1100-4GLTENA ○ ISR1100-4GLTEGB ○ ISR1100-6G • Cisco ISR 1100X Series Integrated Services Routers ○ ISR1100X-4G ○ ISR1100X-6G • Cisco ISR 1000 Series Integrated Series Routers ○ C1101-4P ○ C1101-4PLTEPWX ○ C1101-4PLTEP ○ C1109-2PLTEGB ○ C1109-2PLTEUS ○ C1109-2PLTEVZ ○ C1109-4PLTE2P ○ C1109-4PLTE2PWZ ○ C1111-4P ○ C1111-8P

Claims	Product Categories
	○ C1111-4PLTEEA ○ C1111-4PLTELA ○ C1111-4PW ○ C1111-8PW ○ C1111-8PLTEEA ○ C1111-8PLTELA ○ C1111X-8P ○ C1111-8PLTEEAWA ○ C1111-8PLTEEAWB ○ C1111-8PLTEEAWE ○ C1111-8PLTEEAWR ○ C1111-8PLTEA ○ C1111-8PLTELAWD ○ C1111-8PLTELAWE ○ C1111-8PLTELAWF ○ C1111-8PLTELAWH ○ C1111-8PLTELAWN ○ C1111-8PLTELAWQ ○ C1111-8PLTELAWS ○ C1111-8PLTELAWZ ○ C1111-8PLTELAWA ○ C1111-8PLTEAWY ○ C1111-4PWA ○ C1111-4PWB ○ C1111-4PWD ○ C1111-4PWE ○ C1111-4PWF ○ C1111-4PWH ○ C1111-4PWN ○ C1111-4PWQ ○ C1111-4PWR ○ C1112-8P ○ C1112-8PLTEEA ○ C1112-8PLTEEAWE ○ C1112-8PWE ○ C1113-8PLTEEA ○ C1113-8PLTEEAW ○ C1113-8PLTELA ○ C1113-8PLTELAWZ ○ C1113-8PM ○ C1113-8PMLTEEA ○ C1113-8PMWE ○ C1113-8PW ○ C1113-8PWA

Claims	Product Categories
	○ C1113-8PWB ○ C1113-8PWE ○ C1113-8PWZ ○ C1113-8PLTEAWA ○ C1113-8PLTEEAWA ○ C1113-8PLTEEAWB ○ C1113-8PLTEEAWE ○ C1116-4P ○ C1116-4PLTEEA ○ C1116-4PLTEEAWE ○ C1116-4PLTEEAWA ○ C1116-4PWE ○ C1117-4P ○ C1117-4PLTEEAW ○ C1117-4PMLTEEA ○ C117-4PLTELA ○ C1117-4PLTELAWZ ○ C1117-4PM ○ C1118-8P ○ C1121-4P ○ C1121-4PLTEP ○ C1121-8P ○ C1121-8PLTEP ○ C1121X-8PLTEP ○ C1121X-8PLTEPW ○ C1126-8PLTEP ○ C1127-8PLTEP ○ C1127-8PMLTEP ○ C1128-8PLTEP ○ C1131-8PW ○ C1131-8PLTEPW ○ C1131X-8PW ○ C1131X-8PLTEPW ○ C1161-8P ○ C1161-8PLTEP ○ C1161X-8P ○ C1116-4P ○ C1116-4PLTEEA ○ C1116-4PLTEEAWE ○ C1116-4PLTEEAWA ○ C1116-4PWE ○ C1117-4P ○ C1117-4PLTEEAW ○ C1117-4PMLTEEA

Claims	Product Categories
	○ C117-4PLTELA ○ C1117-4PLTELAWZ ○ C1117-4PM ○ C1118-8P ○ C1121-4P ○ C1121-4PLTEP ○ C1121-8P ○ C1121-8PLTEP ○ C1121X-8PLTEP ○ C1121X-8PLTEPW ○ C1126-8PLTEP ○ C1127-8PLTEP ○ C1127-8PMLTEP ○ C1128-8PLTEP ○ C1131-8PW ○ C1131-8PLTEPW ○ C1131X-8PW ○ C1131X-8PLTEPW ○ C1161-8P ○ C1161-8PLTEP ○ C1161X-8P ○ C1161X-8PLTEP • Cisco Catalyst IR1101 Integrated Services Router Rugged ○ IR-1101-K9 ○ IR-1101-A-K9 • Cisco Catalyst IR1800 Rugged Series Routers ○ IR1821-K9 ○ IR1831-K9 ○ IR1833-K9 ○ IR1835-K9 • Cisco Catalyst IR8100 Heavy Duty Series Routers ○ IR8140H ○ IR8140H-P • Cisco IR8300 Integrated Services Router Rugged ○ IR8340-K9 • Cisco 5000 Series Enterprise Network Compute System ○ ENCS 5104 ○ ENCS 5406 ○ ENCS 5408

Claims	Product Categories
	○ ENCS 5412 (with T1/E1 and 4G NIM modules) • Cisco ESR6300 Embedded Series Routers ○ ESR-6300-NCP-K9 ○ ESR-6300-CON-K9 • Cisco vEdge Devices ○ vEdge 100 ○ vEdge 100b ○ vEdge 100m ○ vEdge 100wm ○ vEdge 1000 ○ vEdge 2000 ○ vEdge 5000
'904 Patent – Claims 1-26	• Cisco Catalyst 2960 Series Switches ○ Catalyst 2960X ○ Catalyst 2960XR • Cisco Catalyst 3750 Series Switches ○ Catalyst 3750 ○ Catalyst 3750X ○ Catalyst 3750G • Cisco Catalyst 9300 Series Switches ○ Catalyst 9300 ○ Catalyst 9300 High Performance ○ Catalyst 9300L ○ Catalyst 9300LM ○ Catalyst 9300X ○ Catalyst 9300 UPOE+ ○ Catalyst 9300 1G • Cisco Catalyst StackWise Platform • Cisco Catalyst StackWise Platform compatible products: ○ Catalyst 9200 Series Switches ○ Catalyst 9400 Series Switches ○ Catalyst 9500 Series Switches ○ Any other Cisco product compatible with the Catalyst StackWise Platform. • Cisco Satellite Network Visualization (nV) System ○ ASR 9000 Series: ▪ Cisco ASR 9001

Claims	Product Categories
	▪ Cisco ASR 9904 ▪ Cisco ASR 9006 ▪ Cisco ASR 9906 ▪ Cisco ASR 9910 ▪ Cisco ASR 9912 ▪ Cisco ASR 9922 ○ Any other Cisco product compatible with the Network Visualization (nV) System. • Cisco 550X Series Stackable Managed Switches ○ SF550X ○ SG550X ○ SX550X
'904 Patent – Claims 4-26	• Cisco 350X Series Stackable Managed Switches ○ SG350X ○ SX350X

EXHIBIT 6

End-of-Sale and End-of-Life Products

These products are no longer being sold.
Click on the product link, when available, for more information.

Please see the **End-of-Life Policy** for more details.

Analytics and Automation Software

Cisco ServiceGrid

Cisco Interfaces and Modules

Cisco 10GBASE Modules

Cisco 10GBASE-CX4 X2 Module
Cisco 10GBASE-ER X2 Module
Cisco 10GBASE-LR X2 Module
Cisco 10GBASE-SR X2 Module

Cisco Access Point Modules

Cisco Aironet Access Point Module for 802.11ac

Cisco ASR 900 Interface Modules

Cisco ASR 900 Series 1-Port 100GE CPAK Module

Cisco Broadband Processing Engines

Cisco UBR-MC20X20V DOCSIS 3.0 Broadband Processing Engine

Cisco uBR-MC88V Broadband Processing Engine

Cisco uBR-MC3GX60V Broadband Processing Engine

Cisco uBR-MC3GX60V DOCSIS3.0 M-CMTS Broadband Processing Engine
Cisco uBR-MC3Gx60V-RPHY Broadband Processing Engine

Cisco Fan Modules

Cisco Catalyst 6807-XL Fan Tray

Cisco Network Modules

Cisco ASR 1000 Series 20-Gbps Embedded Services Processor

Cisco ASR 1000 Series 10Gbps Embedded Services Processor

Cisco ASR 1000 Series 10Gbps Embedded Services Processor Non Crypto

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 450 of 496

Cisco ASR 1000 Series 5Gbps Embedded Services Processor

Cisco Catalyst 6800 Series 8-Port 40 Gigabit Ethernet Module

Cisco Catalyst 6800 Series Supervisor Engine 6T

Cisco Catalyst 6500 Serial 1550nm 10 Gigabit Ethernet Module

Cisco Catalyst 6500 Series 10GBASE-LR Serial 1310nm 10GbE Module

Cisco Catalyst 6500 Series 8-Port Gigabit Ethernet Module

Cisco Catalyst 6500 Series Supervisor Engine 2T

Cisco Catalyst 4500 Supervisor Engine 7-E

Cisco Catalyst 4500 Supervisor Engine 7L-E
Cisco Catalyst C3KX-NM-10G Network Module
Cisco Catalyst C3KX-NM-10GT Network Module
Cisco Catalyst C3KX-NM-1G Network Module
Cisco Nexus 7000 F2-Series 48-Port 1 and 10 Gigabit Ethernet Module
Cisco Nexus 7000 F2-Series Copper 1G and 10G Ethernet Module Enhanced
Cisco Nexus 7000 M2-Series 24-Port 10 Gigabit Ethernet Module
Cisco Nexus 7000 M2-Series 6-Port 40 Gigabit Ethernet Module
Cisco Nexus 7000 M1-Series 48-Port Copper GE Module with XL
Cisco Nexus 7000 M1-Series 48-Port Fiber GE Module with XL
Cisco Nexus 7000 M1-Series 32-Port 10 Gigabit Ethernet Module with XL
Cisco Nexus 7000 M1-Series 8-Port 10 Gigabit Ethernet Module with XL
Cisco Nexus 7000 Series Supervisor 2 Module
Cisco Nexus 7000 Series Supervisor 2E Module

Cisco Network Processing Engines

Cisco uBR7200 Series NPE-G2 Network Processing Engine

Cisco Port Adapters

Cisco 100VG Port Adapter

Cisco Power Supply

Cisco Catalyst 6807-XL Power Converter
Cisco Catalyst 6807-XLPower Supply
Cisco Catalyst 6500 Series 8700W Enhanced AC Power Supply

Cisco Route Processors and Route Switch Processors

Cisco ASR 1000 Series Route Processor (RP1)

Cisco ASR 903 Route Switch Processor 1 (RSP1)
Cisco cBR Series CCAP 160G Supervisor
Cisco uBR10012 Performance Routing Engine 5

Cisco Services Modules

Cisco Catalyst 6500 Series Network Analysis Module (NAM-3)
Cisco Nexus 7000 Series Network Analysis Module (NAM-NX1)

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 451 of 496

Cisco Virtual Security Gateway for Nexus 1000V Series Switch

Cisco Wireless Services Module 2 (WiSM2)

Cisco Shared Port Adapters/SPA Interface Processors

Cisco ASR 1000 Series 10Gbps SPA Interface Processor

Cisco Small Business Network Accessories

Cisco MGBBX1 Gigabit Ethernet BX Mini-GBC SFP Transceiver

Cisco Storage Networking Modules

Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module

Cisco Unified Computing System Adapters

Cisco UCS Virtual Interface Card 1285

Cisco UCS Virtual Interface Card 1280

Cisco UCS Virtual Interface Card 1240

Cisco UCS Virtual Interface Card 1227

Cisco UCS Virtual Interface Card 1227T

Cisco UCS Virtual Interface Card 1225

Cisco Voice Modules and Interface Cards

Cisco Unity Express Integrated Service Engine

Cisco WDM Transmission Modules

Cisco ONS 15454 Optical Booster Amplifier Card

Cisco ONS 15454 Optical Pre-amplifier Card

Cloud and Systems Management

Cisco Application Policy Infrastructure Controller Enterprise Module

Cisco CloudCenter

Cisco CloudCenter Suite

Cisco Crosswork Network Automation

Cisco Crosswork Situation Manager

Cisco Elastic Services Controller

Cisco Elastic Services Controller 4.4
Cisco Elastic Services Controller 4.3
Cisco Elastic Services Controller 4.2
Cisco Elastic Services Controller 4.1
Cisco Elastic Services Controller 4.0
Cisco Elastic Services Controller 3.1

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 452 of 496

- Cisco Elastic Services Controller 3.0
- Cisco Elastic Services Controller 2.3

Cisco Evolved Programmable Network Manager

- Cisco Evolved Programmable Network Manager 4.1
- Cisco Evolved Programmable Network Manager 4.0
- Cisco Evolved Programmable Network Manager 3.1
- Cisco Evolved Programmable Network Manager 3.0
- Cisco Evolved Programmable Network Manager 2.2
- Cisco Evolved Programmable Network Manager 2.1
- Cisco Evolved Programmable Network Manager 2.0

Cisco Intelligent Automation for Cloud

- Cisco Intelligent Automation for Cloud 4.3
- Cisco Intelligent Automation for Cloud 4.3.2
- Cisco Intelligent Automation for Cloud 4.3.1
- Cisco Intelligent Automation for Cloud 4.2
- Cisco Intelligent Automation for Cloud 3.0
- Cisco Network Services Manager for Intelligent Automation for Cloud

Cisco Kinetic

- Cisco Kinetic Data Control Module

Cisco Kinetic for Cities

Cisco NAM 2000 Series Appliances

- Cisco NAM 2404 Appliance
- Cisco Prime NAM 2440 Appliance
- Cisco Prime NAM 2420 Appliance

Cisco NetFlow Generation 3000 Series Appliances

- Cisco NetFlow Generation Appliance (NGA) 3340

Cisco Network Services Orchestrator

- Cisco Network Services Orchestrator 4.6
- Cisco Network Services Orchestrator 4.5
- Cisco Network Services Orchestrator 4.4

Cisco Nexus Dashboard Fabric Controller (Formerly DCNM)

- Cisco Prime Data Center Network Manager 7.2

Cisco Nexus Fabric Manager

Cisco Prime Access Registrar

- Cisco Prime Access Registrar 7.3
- Cisco Prime Access Registrar 7.2
- Cisco Prime Access Registrar 7.1
- Cisco Prime Access Registrar 7.0

Cisco Prime Cable Provisioning

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 453 of 496

Cisco Prime Cable Provisioning 5.3
Cisco Prime Cable Provisioning 5.2
Cisco Prime Cable Provisioning 5.1
Cisco Prime Cable Provisioning 5.0

Cisco Prime Central

Cisco Prime Central 2.1
Cisco Prime Central 2.0
Cisco Prime Central 1.5
Cisco Prime Central 1.5.3
Cisco Prime Central 1.5.2
Cisco Prime Central 1.5.1
Cisco Prime Central 1.4.1

Cisco Prime Fulfillment Multivendor Service Orchestration

Cisco Prime Fulfillment Multivendor Service Orchestration 1.1
Cisco Prime Fulfillment Multivendor Service Orchestration 1.0

Cisco Prime Home

Cisco Prime Home 6.6
Cisco Prime Home 6.5
Cisco Prime Home 6.4
Cisco Prime Home 6.3
Cisco Prime Home 5.2
Cisco Prime Home 5.1
Cisco Prime Home 5.0
Cisco Prime Home 5.X
Cisco Prime Home 3.0
Cisco Prime Home 2.4

Cisco Prime Infrastructure

Cisco Prime Infrastructure 3.2
Cisco Prime Infrastructure 3.1
Cisco Prime Infrastructure 3.0
Cisco Prime Infrastructure 2.2
Cisco Prime Infrastructure 2.1
Cisco Prime Infrastructure 2.0

Cisco Prime IP Express

Cisco Prime IP Express 9.1
Cisco Prime IP Express 9.0
Cisco Prime IP Express 8.3

Cisco Prime IP Express Jumpstart

Cisco Prime IP Express Jumpstart 9.1

Cisco Prime Network

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 454 of 496

- Cisco Prime Network 5.3
- Cisco Prime Network 5.2
- Cisco Prime Network 5.1
- Cisco Prime Network 5.0
- Cisco Prime Network 4.3
- Cisco Prime Network 4.3.2
- Cisco Prime Network 4.3.1
- Cisco Prime Network 4.2.2
- Prime Network 4.2.3

Cisco Prime Network Analysis Module Software

- Cisco Prime Network Analysis Module Software 6.1

Cisco Prime Network Registrar

- Cisco Prime Network Registrar 9.1
- Cisco Prime Network Registrar 9.0
- Cisco Prime Network Registrar 8.3

Cisco Prime Network Services Controller

- Cisco Prime Network Services Controller 3.5
- Cisco Prime Network Services Controller 3.4
- Cisco Prime Network Services Controller 3.3
- Cisco Prime Network Services Controller 3.2
- Cisco Prime Network Services Controller 3.0
- Cisco Prime Network Services Controller 2.1

Cisco Prime Optical

- Cisco Prime Optical 10.7
- Cisco Prime Optical 10.6
- Cisco Prime Optical 10.5

Cisco Prime Performance Manager

- Cisco Prime Performance Manager 1.7

Cisco Prime Provisioning

- Cisco Prime Provisioning 7.2
- Cisco Prime Provisioning 7.1
- Cisco Prime Provisioning 7.0
- Cisco Prime Provisioning 6.8

Cisco Prime Service Catalog

- Cisco Prime Service Catalog 12.1
- Cisco Prime Service Catalog 12.0
- Cisco Prime Service Catalog 11.1

Cisco Process Orchestrator

- Cisco Process Orchestrator 3.5
- Cisco Process Orchestrator 3.5.1

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 455 of 496

Cisco Process Orchestrator 3.4
Cisco Process Orchestrator 3.3
Cisco Process Orchestrator 3.2
Cisco Tidal Ent. Orchestrator Adapter Content Pack for Windows Server
Cisco Tidal Enterprise Orchestrator Adapter for Windows Server

Cisco Transport Manager

Cisco Virtual Topology System

Cisco Virtual Topology System 2.5
Cisco Virtual Topology System 2.4
Cisco Virtual Topology System 2.3
Cisco Virtual Topology System 2.2
Cisco Virtual Topology System 2.1
Cisco Virtual Topology System 2.0

Cisco Virtualized Infrastructure Manager

Cisco Virtualized Infrastructure Manager 2.0

Collaboration Endpoints

Cisco Board Series

Cisco Webex Board 85S
Cisco Webex Board 70S
Cisco Webex Board 55S

Cisco Desk Series

Cisco DX650

Cisco DX70

Cisco Desk Hub
Cisco Webex DX80
Cisco Webex Desk Limited Edition

Cisco IP Communicator

Cisco IP Phone 8800 Series

Cisco Unified IP Conference Phone 8831

Cisco Microphones

Cisco Table Microphone 60
Cisco Table Microphone 20
Cisco TelePresence Audio Science Ceiling Microphone

Cisco Small Business SPA500 Series IP Phones

Cisco SPA514G 4-Line GigE IP Phone

Cisco Small Business Voice Accessories

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 456 of 496

Cisco Wireless-N Bridge for Phone Adapters

Cisco TelePresence IX5000 Series

Cisco TelePresence IX5200
Cisco TelePresence IX5000

Cisco TelePresence MX Series

Cisco TelePresence MX800
Cisco TelePresence MX700
Cisco TelePresence MX300 G2
Cisco TelePresence MX200 G2

Cisco TelePresence Remote Control

Cisco TelePresence Remote Control 5

Cisco TelePresence SX Series

Cisco TelePresence SX80 Codec
Cisco TelePresence SX20 Quick Set
Cisco TelePresence SX10 Quick Set

Cisco TelePresence System EX Series

Cisco TelePresence System EX90

Cisco Unified IP Phone 9900 and 8900 Series Accessories

Cisco Unified IP Color Key Expansion Module
Cisco Unified Video Camera

Cisco Unified IP Phone 8900 Series

Cisco Unified IP Phone 8945

Cisco Unified IP Phone 7900 Series

Cisco Unified IP Phone 7975G
Cisco Unified IP Phone 7965G
Cisco Unified IP Phone 7945G
Cisco Unified IP Phone Expansion Module 7916
Cisco Unified Wireless IP Phone 7926G
Cisco Unified Wireless IP Phone 7925G
Cisco Unified Wireless IP Phone 7920 Multi-Charger

Cisco Unified IP Phones 9900 Series

Cisco Unified IP Phone 9971
Cisco Unified IP Phone 9951

Cisco Webex Share

Cisco Webex Share Device

Tandberg 7000 MXP Dual
Tandberg 2000 MXP

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 457 of 496

Conferencing

Cisco Meeting Server

Acano X-series

Cisco TelePresence Content Server

Cisco TelePresence Management Suite Extensions

Cisco TelePresence Management Suite Provisioning Extension

Cisco TelePresence MCU 5300 Series

Cisco TelePresence MCU 5320

Cisco TelePresence MCU 5310

Cisco TelePresence MCU MSE Series

Cisco TelePresence MCU MSE 8510

Cisco TelePresence MSE 8000 Series

Cisco TelePresence MSE 8000

Cisco TelePresence Supervisor MSE 8050

Cisco WebEx Connect IM

Cisco Webex Meetings Server

Cisco WebEx Meetings Server 4.0

Cisco WebEx Meetings Server 3.0

Cisco WebEx Meetings Server 2.8

Cisco WebEx Meetings Server 2.7

Cisco WebEx Meetings Server 2.6

Cisco WebEx Meetings Server 2.5

Cisco WebEx Meetings Server 2.0

Connected Safety and Security

Cisco IP Camera Applications and Utilities

Cisco Physical Access Gateways

Cisco Physical Access Gateway

Cisco Physical Access Manager

Cisco Video Analytics

Cisco Video Surveillance 8000 Series IP Cameras

Cisco Video Surveillance 8930 IP Camera

Cisco Video Surveillance 8630 IP Camera

Cisco Video Surveillance 8620 IP Camera

Cisco Video Surveillance 8400 IP Camera

Cisco Video Surveillance 8070 IP Camera

Cisco Video Surveillance 8030 IP Camera

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 458 of 496

Cisco Video Surveillance 8020 IP Cameras
Cisco Video Surveillance 8000P IP Cameras

Cisco Video Surveillance 7000 Series IP Cameras

Cisco Video Surveillance 7530PD IP Camera
Cisco Video Surveillance 7070 IP Camera
Cisco Video Surveillance 7030E IP Camera

Cisco Video Surveillance 6000 Series IP Cameras

Cisco Video Surveillance 6630 IP Camera
Cisco Video Surveillance 6620 IP Camera
Cisco Video Surveillance 6500PD IP Camera
Cisco Video Surveillance 6400 IP Camera
Cisco Video Surveillance 6400E IP Camera
Cisco Video Surveillance 6050 IP Camera
Cisco Video Surveillance 6030 IP Camera
Cisco Video Surveillance 6020 IP Camera
Cisco Video Surveillance 6000P IP Camera

Cisco Video Surveillance 3000 Series IP Cameras

Cisco Video Surveillance 3630 IP Camera
Cisco Video Surveillance 3620 IP Camera
Cisco Video Surveillance 3535 IP Camera
Cisco Video Surveillance 3520 IP Camera
Cisco Video Surveillance 3050 IP Camera

Cisco Video Surveillance Encoders

Cisco Video Surveillance 8 Port Encoder
Cisco Video Surveillance 4 Port Encoder

Cisco Video Surveillance Manager
Cisco Video Surveillance Operations Manager Software
Cisco Video Surveillance PTZ IP Cameras

Cisco Video Surveillance 6930 IP Camera
Cisco Video Surveillance 2835 IP Camera
Cisco Video Surveillance 2830 IP Camera

Cisco Video Surveillance Storage System

Cisco Physical Security 4RU Storage Series

Cisco Video Surveillance Virtual Matrix Software
Virtualized Applications for UCS

Contact Center

Cisco Computer Telephony Integration Option

Cisco Computer Telephony Integration Option 11.0(1)
Cisco Computer Telephony Integration Option 10.5

Cisco Enterprise Chat and Email

Cisco Enterprise Chat and Email 11.6(1)

Cisco Finesse

Cisco Finesse 11.6(1)
Cisco Finesse 11.5(1)
Cisco Finesse 11.0(1)
Cisco Finesse 10.5(1)

Cisco Packaged Contact Center Enterprise

Cisco Packaged Contact Center 11.6(2)

Cisco Packaged Contact Center Enterprise 11.6(1)
Cisco Packaged Contact Center Enterprise 11.5(1)
Cisco Packaged Contact Center Enterprise 11.0(3)
Cisco Packaged Contact Center Enterprise 11.0(2)
Cisco Packaged Contact Center Enterprise 11.0(1)
Cisco Packaged Contact Center Enterprise 10.5(1)

Cisco Remote Expert Mobile

Cisco Remote Expert Mobile 11.6(1)
Cisco Remote Expert Mobile 11.5(1)
Cisco Remote Expert Mobile 10.6(3)
Cisco Remote Expert Mobile 10.6(1)

Cisco SocialMiner

Cisco SocialMiner 11.6(1)
Cisco SocialMiner 11.5(1)
Cisco SocialMiner 11.0(1)
Cisco SocialMiner 10.5(1)

Cisco Unified Contact Center Enterprise

Cisco Unified Contact Center Enterprise 11.6(2)

Cisco Unified Contact Center Enterprise 11.6(1)
Cisco Unified Contact Center Enterprise 11.5(1)
Cisco Unified Contact Center Enterprise 11.0(3)
Cisco Unified Contact Center Enterprise 11.0(2)
Cisco Unified Contact Center Enterprise 11.0(1)
Cisco Unified Contact Center Enterprise 10.5(3)
Cisco Unified Contact Center Enterprise 10.5(2)
Cisco Unified Contact Center Enterprise 10.5(1)

Cisco Unified Contact Center Express

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 460 of 496

Cisco Unified Contact Center Express 12.0(1)
Cisco Unified Contact Center Express 11.6(2)
Cisco Unified Contact Center Express 11.6(1)
Cisco Unified Contact Center Express 11.5(1)
Cisco Unified Contact Center Express 11.0(1)
Cisco Unified Contact Center Express 10.6(1)
Cisco Unified Contact Center Express 10.0(1)

Cisco Unified Contact Center Management Portal

Cisco Unified Contact Center Management Portal 11.6(1)
Cisco Unified Contact Center Management Portal 11.5(1)
Cisco Unified Contact Center Management Portal 11.0(1)
Cisco Unified Contact Center Management Portal 10.5(1)

Cisco Unified Customer Voice Portal

Cisco Unified Customer Voice Portal 11.6(1)
Cisco Unified Customer Voice Portal 11.5(1)
Cisco Unified Customer Voice Portal 11.0(1)
Cisco Unified Customer Voice Portal 10.5(1)

Cisco Unified Intelligence Center

Cisco Unified Intelligence Center 11.6(1)
Cisco Unified Intelligence Center 11.5(1)
Cisco Unified Intelligence Center 11.0(1)
Cisco Unified Intelligence Center 10.5(1)
Cisco Unified Intelligence Suite 7.5

Cisco Unified IP Interactive Voice Response (IVR)

Cisco Unified IP Interactive Voice Response (IVR) 12.0(1)
Cisco Unified IP Interactive Voice Response (IVR) 11.6(2)
Cisco Unified IP Interactive Voice Response (IVR) 11.6(1)
Cisco Unified IP Interactive Voice Response (IVR) 11.5(1)
Cisco Unified IP Interactive Voice Response (IVR) 11.0(1)
Cisco Unified IP Interactive Voice Response (IVR) 10.6(1)
Cisco Unified IP Interactive Voice Response (IVR) 10.5(1)
Cisco Unified IP Interactive Voice Response (IVR) 10.0(1)

Cisco Unified Workforce Optimization

Cisco Unified Workforce Optimization Adv Quality Management 2.6(2)
Cisco Unified Workforce Optimization Adv Quality Management 2.6(1)
Cisco Unified Workforce Optimization Adv Quality Management 2.4(1)
Cisco Unified Workforce Optimization Call Recording 2.6(2)
Cisco Unified Workforce Optimization Call Recording 2.6(1)
Cisco Unified Workforce Optimization Call Recording 2.4(1)
Cisco Unified Workforce Optimization Quality Management 2.6(2)

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 461 of 496

Cisco Unified Workforce Optimization Quality Management 2.6(1)
Cisco Unified Workforce Optimization Quality Management 2.4(1)

Cisco Virtualized Voice Browser

Cisco Virtualized Voice Browser 11.6 (1)

Cisco Webex Experience Management (formerly CloudCherry)

Data Center Analytics

Cisco Network Assurance Engine
Cisco Network Insights for Data Center

Cisco Network Insights Advisor
Cisco Network Insights for Resources

Hyperconverged Infrastructure

Cisco HyperFlex HX-Series

Cisco HyperFlex HX240c M4 All Flash Node

Cisco HyperFlex HX240c M4 Node
Cisco HyperFlex HX220c M4 All Flash Node

Cisco HyperFlex HX220c M4 Node

Networking Software (IOS & NX-OS)

Cisco IOS XE 17

Cisco IOS XE Amsterdam 17.3.2
Cisco IOS XE Amsterdam 17.3.1
Cisco IOS XE Amsterdam 17.2.1
Cisco IOS XE Amsterdam 17.1.1
Cisco IOS XE Bengaluru 17.5.1
Cisco IOS XE Bengaluru 17.4.1

Cisco IOS XE 16

Cisco IOS XE Denali 16.3.1
Cisco IOS XE Denali 16.2.1
Cisco IOS XE Denali 16.1.1
Cisco IOS XE Everest 16.6.1
Cisco IOS XE Everest 16.5.1
Cisco IOS XE Everest 16.4.1
Cisco IOS XE Fuji 16.9.1
Cisco IOS XE Fuji 16.8.1
Cisco IOS XE Fuji 16.7.1

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 462 of 496

Cisco IOS XE Gibraltar 16.12.1
Cisco IOS XE Gibraltar 16.11.1
Cisco IOS XE Gibraltar 16.10.1

Cisco IOS XR Software (End-of-Sale)

Cisco IOS XR Software Release 6.7
Cisco IOS XR Software Release 6.6
Cisco IOS XR Software Release 6.5
Cisco IOS XR Software Release 6.4
Cisco IOS XR Software Release 6.3
Cisco IOS XR Software Release 6.2
Cisco IOS XR Software Release 6.1
Cisco IOS XR Software Release 6.0
Cisco IOS XR Software Release 5.3
Cisco IOS XR Software Release 5.2

Optical Networking

Cisco Carrier Packet Transport (CPT) System

Cisco Carrier Packet Transport (CPT) 600
Cisco Carrier Packet Transport (CPT) 50

Cisco ONS 15454 Series Multiservice Transport Platforms

Cisco ONS 15454 M6 Multiservice Transport Platform (MSTP)
Cisco ONS 15454 M2 Multiservice Transport Platform (MSTP)

Routers

Cisco 12000 Series Routers

Cisco 12816 Router
Cisco 12810 Router
Cisco 12416 Router
Cisco 12410 Router
Cisco 12406 Router
Cisco 12404 Router
Cisco 12016 Router
Cisco 12010 Router
Cisco 12006 Router
Cisco 12004 Router

Cisco 7300 Series Routers

Cisco 7301 Router

Cisco 7200 Series Routers

Cisco 7206VXR Router
Cisco 7204VXR Router
Cisco 7201 Router

Cisco 5900 Series Embedded Services Routers

Cisco 5940 Embedded Services Router

Cisco 5915 Embedded Service Router

Cisco 3900 Series Integrated Services Routers

Cisco 3945 Integrated Services Router
Cisco 3945E Integrated Services Router
Cisco 3925 Integrated Services Router
Cisco 3925E Integrated Services Router

Cisco 2900 Series Integrated Services Routers

Cisco 2951 Integrated Services Router
Cisco 2921 Integrated Services Router
Cisco 2911 Integrated Services Router
Cisco 2911A Integrated Services Router
Cisco 2901 Integrated Services Router

Cisco 1900 Series Integrated Services Routers

Cisco 1981 Integrated Services Router
Cisco 1941 Integrated Services Router
Cisco 1941W Integrated Services Router
Cisco 1921 Integrated Services Router
Cisco 1905 Serial Integrated Services Router

Cisco 900 Series Industrial Routers

Cisco 910 Industrial Router

Cisco 800 Series Industrial Integrated Services Routers

Cisco 829 Industrial Integrated Services Routers

Cisco 809 Industrial Integrated Services Routers

Cisco 807 Industrial Integrated Services Routers

Cisco 800 Series Routers

Cisco C899 Secure Gigabit Ethernet with Multi-mode 4G LTE Router
Cisco C898EA Integrated Services Router
Cisco C887VAM Integrated Series Routers
Cisco C886VA Integrated Services Routers
Cisco C886VAJ Integrated Services Routers

Cisco ASR 9000 Series Aggregation Services Routers

Cisco ASR 9001 Router

Cisco ASR 1000 Series Aggregation Services Routers

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 464 of 496

- Cisco ASR 1013 Router
- Cisco ASR 1002 Router
- Cisco ASR 1002-X Router
- Cisco ASR 1001 Router
- Cisco ASR 1001-X Router

Cisco ASR 901 Series Aggregation Services Routers

- Cisco ASR 901-12C-F-D Router
- Cisco ASR 901-12C-FT-D Router
- Cisco ASR 901-4C-F-D Router
- Cisco ASR 901-4C-FT-D Router

Cisco ASR 901S Series Aggregation Services Routers

- Cisco ASR 901S-4SG-F-D Router
- Cisco ASR 901S-3SG-F-AH Router
- Cisco ASR 901S-3SG-F-D Router
- Cisco ASR 901S-2SG-F-AH Router
- Cisco ASR 901S-2SG-F-D Router

Cisco Integrated Services Virtual Router

Cisco Network Convergence System 5500 Series

- Cisco Network Convergence System 5502
- Cisco Network Convergence System 5502-SE

Cisco Network Convergence System 5000 Series

- Cisco Network Convergence System 5011

Cisco Small Business RV Series Routers

- Cisco CVR100W Wireless-N VPN Router
- Cisco RV345 Dual WAN Gigabit VPN Router
- Cisco RV345P Dual WAN Gigabit POE VPN Router
- Cisco RV340 Dual WAN Gigabit VPN Router
- Cisco RV340W Dual WAN Gigabit Wireless-AC VPN Router
- Cisco RV325 Dual Gigabit WAN VPN Router
- Cisco RV325 Dual Gigabit WAN WF VPN Router
- Cisco RV320 Dual Gigabit WAN VPN Router
- Cisco RV320 Dual Gigabit WAN WF VPN Router
- Cisco RV315W Wireless-N VPN Router
- Cisco RV260 VPN Router
- Cisco RV260P VPN Router with PoE
- Cisco RV260W Wireless-AC VPN Router
- Cisco RV215W Wireless-N VPN Router
- Cisco RV160 VPN Router
- Cisco RV160W Wireless-AC VPN Router
- Cisco RV134W VDSL2 Wireless-AC VPN Router
- Cisco RV132W ADSL2+ Wireless-N VPN Router

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 465 of 496

Cisco RV130 WF VPN Router
Cisco RV130W Wireless-N Multifunction VPN Router
Cisco RV130W Wireless-N Multifunction VPN Router WF
Cisco RV110W Wireless-N VPN Firewall
Cisco RV082 Dual WAN VPN Router
Cisco RV042 Dual WAN VPN Router
Cisco RV042G Dual Gigabit WAN VPN Router
Cisco RV016 Multi-WAN VPN Router
Cisco RVL200 4-Port SSL/IPsec VPN Router

Cisco WAN Automation Engine (WAE)

Cisco WAN Automation Engine EoS releases

Cisco Wide Area Virtualization Engines

Cisco WAVE 8541 Wide Area Virtualization Engine
Cisco WAVE 7571 Wide Area Virtualization Engine
Cisco WAVE 7541 Wide Area Virtualization Engine
Cisco WAVE 694 Wide Area Virtualization Engine
Cisco WAVE 594 Wide Area Virtualization Engine
Cisco WAVE 294 Wide Area Virtualization Engine

Cisco XR 12000 Series Router

Cisco XR 12416 Router
Cisco XR 12410 Router
Cisco XR 12406 Router
Cisco XR 12404 Router

Security

Cisco AMP for Networks

Cisco AMP 8150
Cisco AMP 7150

Cisco ASA 5500-X Series Firewalls

Cisco ASA 5585-X Adaptive Security Appliance
Cisco ASA 5585-X with No Payload Encryption
Cisco ASA 5555-X Adaptive Security Appliance
Cisco ASA 5515-X Adaptive Security Appliance
Cisco ASA 5515-X Adaptive Security Appliance - No Payload Encryption
Cisco ASA 5512-X Adaptive Security Appliance
Cisco ASA 5512-X Adaptive Security Appliance - No Payload Encryption
Cisco ASA 5505 Adaptive Security Appliance

Cisco ASA 5500-X with FirePOWER Services

Cisco ASA 5585-X with FirePOWER SSP-60
Cisco ASA 5585-X with FirePOWER SSP-40

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 466 of 496

Cisco ASA 5585-X with FirePOWER SSP-20
Cisco ASA 5585-X with FirePOWER SSP-10
Cisco ASA 5555-X with FirePOWER Services
Cisco ASA 5545-X with FirePOWER Services
Cisco ASA 5525-X with FirePOWER Services
Cisco ASA 5516-X with FirePOWER Services
Cisco ASA 5515-X with FirePOWER Services
Cisco ASA 5512-X with FirePOWER Services
Cisco ASA 5508-X with FirePOWER Services
Cisco ASA 5506-X with FirePOWER Services
Cisco ASA 5506H-X with FirePOWER Services
Cisco ASA 5506W-X with FirePOWER Services

Cisco Centri Firewall

Cisco Centri Firewall Patches

Cisco Compatible IntraGuard Firewall Series

Cisco FirePOWER 8000 Series Appliances

Cisco FirePOWER Appliance 8390
Cisco FirePOWER Appliance 8370
Cisco FirePOWER Appliance 8360
Cisco FirePOWER Appliance 8350
Cisco FirePOWER Appliance 8140
Cisco FirePOWER Appliance 8130
Cisco FirePOWER Appliance 8120

Cisco FirePOWER 7000 Series Appliances

Cisco FirePOWER Appliance 7125
Cisco FirePOWER Appliance 7120
Cisco FirePOWER Appliance 7115
Cisco FirePOWER Appliance 7110
Cisco FirePOWER Appliance 7050
Cisco FirePOWER Appliance 7030
Cisco FirePOWER Appliance 7020
Cisco FirePOWER Appliance 7010

Cisco Firepower 4100 Series

Cisco Firepower 4110 Security Appliance

Cisco Identity Services Engine

Cisco Identity Services Engine 2.6
Cisco Identity Services Engine 2.4
Cisco Identity Services Engine 2.2
Cisco Identity Services Engine 2.0
Cisco Identity Services Engine 1.0.4

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 467 of 496

Cisco IDS Host Sensors
Cisco Incident Control System
Cisco Intrusion Prevention System
Cisco IOS Content Filtering
Cisco IPS 4200 Series Sensors

- Cisco IPS 4240 Sensor

Cisco Multinet
Cisco NAC Legacy Software
Cisco Secure Access Control System

- Cisco Secure Access Control System 5.8
- Cisco Secure Access Control System 5.8.1
- Cisco Secure Access Control System Migration Tool

Cisco Secure Email and Web Manager

- Cisco Content Security Management Appliance M690
- Cisco Content Security Management Appliance M690X
- Cisco Content Security Management Appliance M680
- Cisco Content Security Management Appliance M390
- Cisco Content Security Management Appliance M390X
- Cisco Content Security Management Appliance M380
- Cisco Content Security Management Appliance M190
- Cisco Content Security Management Appliance M170

Cisco Secure Email Gateway

- Cisco Email Security Appliance C690
- Cisco Email Security Appliance C690X
- Cisco Email Security Appliance C680
- Cisco Email Security Appliance C390
- Cisco Email Security Appliance C380
- Cisco Email Security Appliance C190
- Cisco Email Security Appliance C170

Cisco Secure Firewall Management Center

- Cisco FireSIGHT Management Center 750
- Cisco Firepower Management Center 4000
- Cisco Firepower Management Center 2000

Cisco Secure Malware Analytics (Threat Grid)

- Cisco Threat Grid 5504 Appliance
- Cisco Threat Grid 5500 Appliance
- Cisco Threat Grid 5004 Appliance
- Cisco Threat Grid 5000 Appliance

Cisco Secure Network Analytics Flow Collector

- Cisco Stealthwatch Flow Collector 5200
- Cisco Stealthwatch Flow Collector 5020
- Cisco Stealthwatch Flow Collector 5000
- Cisco Stealthwatch Flow Collector 4200
- Cisco Stealthwatch Flow Collector 4010
- Cisco Stealthwatch Flow Collector 2010
- Cisco Stealthwatch Flow Collector 1010

Cisco Secure Network Analytics Flow Sensor

- Cisco Stealthwatch Flow Sensor 4210
- Cisco Stealthwatch Flow Sensor 4200
- Cisco Stealthwatch Flow Sensor 4010
- Cisco Stealthwatch Flow Sensor 3200
- Cisco Stealthwatch Flow Sensor 3010
- Cisco Stealthwatch Flow Sensor 2200
- Cisco Stealthwatch Flow Sensor 2010
- Cisco Stealthwatch Flow Sensor 1200
- Cisco Stealthwatch Flow Sensor 1010

Cisco Secure Network Analytics Manager

- Cisco Stealthwatch Management Console 2200
- Cisco Stealthwatch Management Console 2010
- Cisco Stealthwatch Management Console 2000
- Cisco Stealthwatch Management Console 1010
- Cisco Stealthwatch Management Console 1000

Cisco Secure Network Analytics UDP Director

- Cisco Stealthwatch UDP Director 2200
- Cisco Stealthwatch UDP Director 2010
- Cisco Stealthwatch UDP Director 1010

Cisco Secure Network Server (SNS) 3400 Series

- Cisco Secure Network Server (SNS) 3495 Appliance
- Cisco Secure Network Server (SNS) 3415 Appliance

Cisco Secure Web Appliance

- Cisco Web Security Appliance S690
- Cisco Web Security Appliance S690X
- Cisco Web Security Appliance S680
- Cisco Web Security Appliance S380
- Cisco Web Security Appliance S170

Cisco Security Manager

- Cisco Security Manager 4.15
- Cisco Security Manager 4.14
- Cisco Security Manager 4.13
- Cisco Security Manager 4.12

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 469 of 496

Cisco Security Manager 4.11
Cisco Security Manager 4.10
Cisco Security Manager 4.9

Cisco Security Packet Analyzer

Cisco Security Packet Analyzer 2400

Cisco SSL Appliances

SSL Appliance 8200
SSL Appliance 2000
SSL Appliance 1500

Cisco Stealthwatch Proxy License

Servers - Unified Computing

Cisco C800 Series

Cisco C880 M4 Server

Cisco C880 M4 Storage Subsystem
Cisco C880 M4 with v4 CPUs Storage Subsystem
Cisco C880 M4 with v3 CPUs Server
Cisco C880 M4 with v3 CPUs Storage Subsystem

Cisco R Series Racks

Cisco R42612 Rack
Cisco R42610 Rack

Cisco RP Series Power Distribution Units

Cisco UCS 6200 Series Fabric Interconnects

Cisco UCS 6296UP 96-Port Fabric Interconnect
Cisco UCS 6248UP 48-Port Fabric Interconnect

Cisco UCS B-Series Blade Servers

Cisco UCS B460 M4 Blade Server

Cisco UCS B420 M4 Blade Server

Cisco UCS B420 M3 Blade Server

Cisco UCS B260 M4 Blade Server

Cisco UCS B200 M4 Blade Server

Cisco UCS B200 M3 Blade Server

Cisco UCS C-Series Rack Servers

Cisco UCS C3160 Rack Server

Cisco UCS C480 ML M5 Rack Server

Cisco UCS C460 M4 Rack Server

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 470 of 496

Cisco UCS C420 M3 Rack Server

Cisco UCS C240 M4 Rack Server

Cisco UCS C240 M3 Rack Server

Cisco UCS C220 M4 Rack Server

Cisco UCS C220 M3 Rack Server

Cisco UCS C-Series Rack-Mount Standalone Server Software

Cisco UCS C480 ML M5 Rack Server Software

Cisco UCS M-Series Modular Servers

Cisco UCS M4308 Modular Chassis
Cisco UCS M2814 Compute Cartridge
Cisco UCS M1414 Compute Cartridge
Cisco UCS M142 Compute Cartridge

Storage Networking

Cisco MDS 9500 Series Multilayer Directors

Cisco MDS 9513 Multilayer Director
Cisco MDS 9506 Multilayer Director

Cisco MDS 9200 Series Multiservice Switches

Cisco MDS 9216/9216A Multilayer Fabric Switch
Cisco MDS 9216A Multilayer Fabric Switch
Cisco MDS 9216i Multilayer Fabric Switch

Cisco MDS 9020 Series Fabric Switch

Cisco MDS 9000 Services-Oriented SANs

Cisco Data Mobility Manager

Cisco MDS Blade Switch Series

Cisco MDS 8G FC HP Blade Switch

Switches

Cisco 6000 Series IP DSL Switches

Cisco 2500 Series Connected Grid Switches

Cisco CGS-2520-24TC Connected Grid Switch
Cisco CGS-2520-16S-8PC Connected Grid Switch

Cisco 550X Series Stackable Managed Switches

Cisco SF550X-48 48-Port 10/100 Stackable Managed Switch
Cisco SF550X-48MP 48-Port 10/100 PoE Stackable Managed Switch
Cisco SF550X-48P 48-Port 10/100 PoE Stackable Managed Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 471 of 496

Cisco SF550X-24 24-Port 10/100 Stackable Managed Switch
Cisco SF550X-24MP 24-Port 10/100 PoE Stackable Managed Switch
Cisco SF550X-24P 24-Port 10/100 PoE Stackable Managed Switch
Cisco SG550X-48 48-Port Gigabit Stackable Managed Switch
Cisco SG550X-48MP 48-Port Gigabit PoE Stackable Managed Switch
Cisco SG550X-48P 48-Port Gigabit PoE Stackable Managed Switch
Cisco SG550X-24 24-Port Gigabit Stackable Managed Switch
Cisco SG550X-24MP 24-Port Gigabit PoE Stackable Managed Switch
Cisco SG550X-24MPP 24-Port Gigabit PoE Stackable Managed Switch
Cisco SG550X-24P 24-Port Gigabit PoE Stackable Managed Switch
Cisco SG550XG-48T 48-Port 10GBase-T Stackable Managed Switch
Cisco SG550XG-24F 24-Port 10G SFP+ Stackable Managed Switch
Cisco SG550XG-24T 24-Port 10GBase-T Stackable Managed Switch
Cisco SG550XG-8F8T 16-Port 10G Stackable Managed Switch
Cisco SX550X-52 52-Port 10GBase-T Stackable Managed Switch
Cisco SX550X-24 24-Port 10GBase-T Stackable Managed Switch
Cisco SX550X-24F 24-Port 10G SFP+ Stackable Managed Switch
Cisco SX550X-24FT 24-Port 10G Stackable Managed Switch
Cisco SX550X-16FT 16-Port 10G Stackable Managed Switch
Cisco SX550X-12F 12-Port 10G SFP+ Stackable Managed Switch

Cisco 350 Series Managed Switches

Cisco SF352-08 8-Port 10/100 Managed Switch
Cisco SF352-08MP 8-Port 10/100 POE Managed Switch
Cisco SF352-08P 8-Port 10/100 POE Managed Switch
Cisco SF350-48 48-Port 10/100 Managed Switch
Cisco SF350-48MP 48-Port 10/100 PoE Managed Switch
Cisco SF350-48P 48-Port 10/100 PoE Managed Switch
Cisco SF350-24 24-Port 10/100 Managed Switch
Cisco SF350-24MP 24-Port 10/100 Max PoE Managed Switch
Cisco SF350-24P 24-Port 10/100 POE Managed Switch
Cisco SF350-08 8-Port 10/100 Managed Switch
Cisco SG355-10P 10-Port Gigabit PoE Managed Switch
Cisco SG350-52 52-Port Gigabit Managed Switch
Cisco SG350-52MP 52-Port Gigabit Max-PoE Managed Switch
Cisco SG350-52P 52-Port Gigabit PoE Managed Switch
Cisco SG350-28 28-Port Gigabit Managed Switch
Cisco SG350-28MP 28-Port Gigabit PoE Managed Switch
Cisco SG350-28P 28-Port Gigabit PoE Managed Switch
Cisco SG350-28SFP 28-Port Gigabit Managed SFP Switch
Cisco SG350-20 20-Port Gigabit Managed Switch
Cisco SG350-10 10-Port Gigabit Managed Switch
Cisco SG350-10MP 10-Port Gigabit PoE Managed Switch
Cisco SG350-10P 10-Port Gigabit PoE Managed Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 472 of 496

Cisco SG350-10SFP 10-Port Gigabit Managed SFP Switch
Cisco SG350-8PD 8-Port 2.5G PoE Managed Switch

Cisco 350X Series Stackable Managed Switches

Cisco SG350X-48 48-Port Gigabit Stackable Managed Switch
Cisco SG350X-48MP 48-Port Gigabit PoE Stackable Managed Switch
Cisco SG350X-48P 48-Port Gigabit PoE Stackable Managed Switch
Cisco SG350X-48PV 48-Port 5G PoE Stackable Managed Switch
Cisco SG350X-24 24-Port Gigabit Stackable Managed Switch
Cisco SG350X-24MP 24-Port Gigabit PoE Stackable Managed Switch
Cisco SG350X-24P 24-Port Gigabit PoE Stackable Managed Switch
Cisco SG350X-24PD 24-Port 2.5G PoE Stackable Managed Switch
Cisco SG350X-24PV 24-Port 5G PoE Stackable Managed Switch
Cisco SG350X-12PMV 12-Port 5G PoE Stackable Managed Switch
Cisco SG350X-8PMD 8-Port 2.5G PoE Stackable Managed Switch
Cisco SG350XG-48T 48-Port 10GBase-T Stackable Managed Switch
Cisco SG350XG-24F 24-Port 10G SFP+ Stackable Managed Switch
Cisco SG350XG-24T 24-Port 10GBase-T Stackable Managed Switch
Cisco SG350XG-2F10 12-Port 10GBase-T Stackable Managed Switch
Cisco SX350X-52 52-Port 10GBase-T Stackable Managed Switch
Cisco SX350X-24 24-Port 10GBase-T Stackable Managed Switch
Cisco SX350X-24F 24-Port 10G SFP+ Stackable Managed Switch
Cisco SX350X-12 12-Port 10GBase-T Stackable Managed Switch
Cisco SX350X-08 8-Port 10GBase-T Stackable Managed Switch

Cisco 250 Series Smart Switches

Cisco SF250-48 48-Port 10/100 Smart Switch
Cisco SF250-48HP 48-Port 10/100 PoE Smart Switch
Cisco SF250-24 24-Port 10/100 Smart Switch
Cisco SF250-24P 24-Port 10/100 PoE Smart Switch
Cisco SG250-50 50-Port Gigabit Smart Switch
Cisco SG250-50HP 50-Port Gigabit PoE Smart Switch
Cisco SG250-50P 50-Port Gigabit PoE Smart Switch
Cisco SG250-26 26-Port Gigabit Smart Switch
Cisco SG250-26HP 26-Port Gigabit PoE Smart Switch
Cisco SG250-26P 26-Port Gigabit PoE Smart Switch
Cisco SG250-18 18-Port Gigabit Smart Switch
Cisco SG250-10P 10-Port Gigabit PoE Smart Switch
Cisco SG250-08 8-Port Gigabit Smart Switch
Cisco SG250-08HP 8-Port Gigabit PoE Smart Switch
Cisco SG250X-48 48-Port Gigabit with 4-Port 10-Gigabit Smart Switch
Cisco SG250X-48P Gigabit PoE with 4-Port 10-Gigabit Smart Switch
Cisco SG250X-24 24-Port Gigabit with 4-Port 10-Gigabit Smart Switch
Cisco SG250X-24P Gigabit PoE with 4-Port 10-Gigabit Smart Switch

Cisco 220 Series Smart Switches

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 473 of 496

Cisco SF220-48 48-Port 10/100 Smart Switch
Cisco SF220-48P 48-Port 10/100 PoE Smart Switch
Cisco SF220-24 24-Port 10/100 Smart Switch
Cisco SF220-24P 24-Port 10/100 PoE Smart Switch
Cisco SG220-52 52-Port Gigabit Smart Switch
Cisco SG220-50 50-Port Gigabit Smart Switch
Cisco SG220-50P 50-Port Gigabit PoE Smart Switch
Cisco SG220-28 28-Port Gigabit Smart Switch
Cisco SG220-28MP 28-Port Gigabit PoE Smart Switch
Cisco SG220-26 26-Port Gigabit Smart Switch
Cisco SG220-26P 26-Port Gigabit PoE Smart Switch

Cisco Blade Switches for HP

Cisco Catalyst Blade Switch 3120 for HP

Cisco Catalyst 6800 Series Switches

Cisco Catalyst 6880-X Switch
Cisco Catalyst 6840-X Switch
Cisco Catalyst 6807-XL Switch
Cisco Catalyst 6800ia Switch

Cisco Catalyst C6840-X-LE-40G Switch
Cisco Catalyst C6832-X-LE Switch
Cisco Catalyst C6824-X-LE-40G Switch
Cisco Catalyst C6816-X-LE Switch

Cisco Catalyst 6500 Series Switches

Cisco Catalyst 6513-E Switch
Cisco Catalyst 6509-E Switch
Cisco Catalyst 6509-NEB-A Switch

Cisco Catalyst 6509-V-E Switch

Cisco Catalyst 6506-E Switch
Cisco Catalyst 6504-E Switch

Cisco Catalyst 6503-E Switch

Cisco Catalyst 4900 Series Switches

Cisco Catalyst 4948E Ethernet Switch

Cisco Catalyst 4900M Switch

Cisco Catalyst 4500 Series Switches

Cisco Catalyst 4510R Switch
Cisco Catalyst 4510R+E Switch
Cisco Catalyst 4507R+E Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 474 of 496

- Cisco Catalyst 4506-E Switch
- Cisco Catalyst 4503-E Switch

Cisco Catalyst 4500-X Series Switches

- Cisco Catalyst 4500X-40 SFP+ Switch
- Cisco Catalyst 4500X-32 SFP+ Switch
- Cisco Catalyst 4500X-24 SFP+ Switch
- Cisco Catalyst 4500X-16 SFP+ Switch
- Cisco Catalyst 4500X-F-32 SFP+ Switch
- Cisco Catalyst 4500X-F-16 SFP+ Switch

Cisco Catalyst 3850 Series Switches

- Cisco Catalyst 3850-48F-E Switch
- Cisco Catalyst 3850-48F-L Switch
- Cisco Catalyst 3850-48F-S Switch
- Cisco Catalyst 3850-48P-E Switch
- Cisco Catalyst 3850-48P-L Switch
- Cisco Catalyst 3850-48P-S Switch
- Cisco Catalyst 3850-48PW-S Bundle
- Cisco Catalyst 3850-48T-E Switch
- Cisco Catalyst 3850-48T-L Switch
- Cisco Catalyst 3850-48T-S Switch
- Cisco Catalyst 3850-48U-E Switch
- Cisco Catalyst 3850-48U-L Switch
- Cisco Catalyst 3850-48U-S Switch
- Cisco Catalyst 3850-48XS-E Switch
- Cisco Catalyst 3850-48XS-F-E Switch
- Cisco Catalyst 3850-48XS-F-S Switch
- Cisco Catalyst 3850-48XS-S Switch
- Cisco Catalyst 3850-32XS-E Switch
- Cisco Catalyst 3850-32XS-S Switch
- Cisco Catalyst 3850-24P-E Switch
- Cisco Catalyst 3850-24P-L Switch
- Cisco Catalyst 3850-24P-S Switch
- Cisco Catalyst 3850-24PW-S Bundle
- Cisco Catalyst 3850-24S-E Switch
- Cisco Catalyst 3850-24S-S Switch
- Cisco Catalyst 3850-24T-E Switch
- Cisco Catalyst 3850-24T-L Switch
- Cisco Catalyst 3850-24T-S Switch
- Cisco Catalyst 3850-24U-E Switch
- Cisco Catalyst 3850-24U-L Switch
- Cisco Catalyst 3850-24U-S Switch
- Cisco Catalyst 3850-24XS-E Switch
- Cisco Catalyst 3850-24XS-S Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 475 of 496

Cisco Catalyst 3850-24XU-E Switch
Cisco Catalyst 3850-24XU-L Switch
Cisco Catalyst 3850-24XU-S Switch
Cisco Catalyst 3850-16XS-E Switch
Cisco Catalyst 3850-16XS-S Switch
Cisco Catalyst 3850-12S-E Switch
Cisco Catalyst 3850-12S-S Switch
Cisco Catalyst 3850-12XS-E Switch
Cisco Catalyst 3850-12XS-S Switch
Cisco Catalyst C3850-12X48U-E Switch
Cisco Catalyst C3850-12X48U-L Switch
Cisco Catalyst C3850-12X48U-S Switch

Cisco Catalyst 3750 Series Switches

Cisco Catalyst 3750V2-48PS Switch
Cisco Catalyst 3750V2-48TS Switch
Cisco Catalyst 3750V2-24FS Switch
Cisco Catalyst 3750V2-24PS Switch
Cisco Catalyst 3750V2-24TS Switch

Cisco Catalyst 3750-X Series Switches

Cisco Catalyst 3750X-48P-E Switch
Cisco Catalyst 3750X-48P-L Switch
Cisco Catalyst 3750X-48P-S Switch
Cisco Catalyst 3750X-48PF-E Switch
Cisco Catalyst 3750X-48PF-L Switch
Cisco Catalyst 3750X-48PF-S Switch
Cisco Catalyst 3750X-48T-E Switch
Cisco Catalyst 3750X-48T-L Switch
Cisco Catalyst 3750X-48T-S Switch
Cisco Catalyst 3750X-48U-E Switch
Cisco Catalyst 3750X-48U-L Switch
Cisco Catalyst 3750X-48U-S Switch
Cisco Catalyst 3750X-24P-E Switch
Cisco Catalyst 3750X-24P-L Switch
Cisco Catalyst 3750X-24P-S Switch
Cisco Catalyst 3750X-24S-E Switch
Cisco Catalyst 3750X-24S-S Switch
Cisco Catalyst 3750X-24T-E Switch
Cisco Catalyst 3750X-24T-L Switch
Cisco Catalyst 3750X-24T-S Switch
Cisco Catalyst 3750X-24U-E Switch
Cisco Catalyst 3750X-24U-L Switch
Cisco Catalyst 3750X-24U-S Switch
Cisco Catalyst 3750X-12S-E Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 476 of 496

Cisco Catalyst 3750X-12S-S Switch
Cisco Catalyst 3K-X Fan Module

Cisco Catalyst 3650 Series Switches

Cisco Catalyst 3650-48FD-E Switch
Cisco Catalyst 3650-48FD-L Switch
Cisco Catalyst 3650-48FD-S Switch
Cisco Catalyst 3650-48FQ-E Switch
Cisco Catalyst 3650-48FQ-L Switch
Cisco Catalyst 3650-48FQ-S Switch
Cisco Catalyst 3650-48FS-E Switch
Cisco Catalyst 3650-48FS-L Switch
Cisco Catalyst 3650-48FS-S Switch
Cisco Catalyst 3650-48PD-E Switch
Cisco Catalyst 3650-48PD-L Switch
Cisco Catalyst 3650-48PD-S Switch
Cisco Catalyst 3650-48PQ-E Switch
Cisco Catalyst 3650-48PQ-L Switch
Cisco Catalyst 3650-48PQ-S Switch
Cisco Catalyst 3650-48PS-E Switch
Cisco Catalyst 3650-48PS-L Switch
Cisco Catalyst 3650-48PS-S Switch
Cisco Catalyst 3650-48TD-E Switch
Cisco Catalyst 3650-48TD-L Switch
Cisco Catalyst 3650-48TD-S Switch
Cisco Catalyst 3650-48TQ-E Switch
Cisco Catalyst 3650-48TQ-L Switch
Cisco Catalyst 3650-48TQ-S Switch
Cisco Catalyst 3650-48TS-E Switch
Cisco Catalyst 3650-48TS-L Switch
Cisco Catalyst 3650-48TS-S Switch
Cisco Catalyst 3650-24PD-E Switch
Cisco Catalyst 3650-24PD-L Switch
Cisco Catalyst 3650-24PD-S Switch
Cisco Catalyst 3650-24PS-E Switch
Cisco Catalyst 3650-24PS-L Switch
Cisco Catalyst 3650-24PS-S Switch
Cisco Catalyst 3650-24TD-E Switch
Cisco Catalyst 3650-24TD-L Switch
Cisco Catalyst 3650-24TD-S Switch
Cisco Catalyst 3650-24TS-E Switch
Cisco Catalyst 3650-24TS-L Switch
Cisco Catalyst 3650-24TS-S Switch
Cisco Catalyst 3650-12X48FD-E Switch
Cisco Catalyst 3650-12X48FD-L Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 477 of 496

- Cisco Catalyst 3650-12X48FD-S Switch
- Cisco Catalyst 3650-12X48UQ-E Switch
- Cisco Catalyst 3650-12X48UQ-L Switch
- Cisco Catalyst 3650-12X48UQ-S Switch
- Cisco Catalyst 3650-12X48UR-E Switch
- Cisco Catalyst 3650-12X48UR-L Switch
- Cisco Catalyst 3650-12X48UR-S Switch
- Cisco Catalyst 3650-12X48UZ-E Switch
- Cisco Catalyst 3650-12X48UZ-L Switch
- Cisco Catalyst 3650-12X48UZ-S Switch
- Cisco Catalyst 3650-8X24PD-E Switch
- Cisco Catalyst 3650-8X24PD-L Switch
- Cisco Catalyst 3650-8X24PD-S Switch
- Cisco Catalyst 3650-8X24UQ-E Switch
- Cisco Catalyst 3650-8X24UQ-L Switch
- Cisco Catalyst 3650-8X24UQ-S Switch

Cisco Catalyst 3560 Series Switches

- Cisco Catalyst 3560V2-48PS Switch
- Cisco Catalyst 3560V2-48TS Switch
- Cisco Catalyst 3560V2-24DC Switch
- Cisco Catalyst 3560V2-24PS Switch
- Cisco Catalyst 3560V2-24TS Switch

Cisco Catalyst 3560-C Series Switches

- Cisco Catalyst 3560C-12PC-s Switch
- Cisco Catalyst 3560C-8PC-S Switch
- Cisco Catalyst 3560CG-8PC-S Compact Switch
- Cisco Catalyst 3560CG-8TC-S Compact Switch
- Cisco Catalyst 3560CPD-8PT-S Compact Switch

Cisco Catalyst 3560-X Series Switches

- Cisco Catalyst 3560X-48P-E Switch
- Cisco Catalyst 3560X-48P-L Switch
- Cisco Catalyst 3560X-48P-S Switch
- Cisco Catalyst 3560X-48PF-E Switch
- Cisco Catalyst 3560X-48PF-L Switch
- Cisco Catalyst 3560X-48PF-S Switch
- Cisco Catalyst 3560X-48T-E Switch
- Cisco Catalyst 3560X-48T-L Switch
- Cisco Catalyst 3560X-48T-S Switch
- Cisco Catalyst 3560X-48U-E Switch
- Cisco Catalyst 3560X-48U-L Switch
- Cisco Catalyst 3560X-48U-S Switch
- Cisco Catalyst 3560X-24P-E Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 478 of 496

- Cisco Catalyst 3560X-24P-L Switch
- Cisco Catalyst 3560X-24P-S Switch
- Cisco Catalyst 3560X-24T-E Switch
- Cisco Catalyst 3560X-24T-L Switch
- Cisco Catalyst 3560X-24T-S Switch
- Cisco Catalyst 3560X-24U-E Switch
- Cisco Catalyst 3560X-24U-L Switch
- Cisco Catalyst 3560X-24U-S Switch

Cisco Catalyst 2960 Series Switches

- Cisco Catalyst 2960-48PST-L Switch
- Cisco Catalyst 2960-48PST-S Switch
- Cisco Catalyst 2960-48TC-L Switch
- Cisco Catalyst 2960-48TC-S Switch
- Cisco Catalyst 2960-48TT-L Switch
- Cisco Catalyst 2960-48TT-S Switch
- Cisco Catalyst 2960-24-S Switch
- Cisco Catalyst 2960-24LC-S Switch
- Cisco Catalyst 2960-24LT-L Switch
- Cisco Catalyst 2960-24PC-L Switch
- Cisco Catalyst 2960-24PC-S Switch
- Cisco Catalyst 2960-24TC-L Switch
- Cisco Catalyst 2960-24TC-S Switch
- Cisco Catalyst 2960-24TT-L Switch

Cisco Catalyst 2960-C Series Switches

- Cisco Catalyst 2960CG-8TC-L Compact Switch

Cisco Catalyst 2960-L Series Switches

- Cisco Catalyst 2960L-48PQ-LL Switch
- Cisco Catalyst 2960L-48PS-LL Switch
- Cisco Catalyst 2960L-48TQ-LL Switch
- Cisco Catalyst 2960L-48TS-LL Switch
- Cisco Catalyst 2960L-24PQ-LL Switch
- Cisco Catalyst 2960L-24PS-LL Switch
- Cisco Catalyst 2960L-24TQ-LL Switch
- Cisco Catalyst 2960L-24TS-LL Switch
- Cisco Catalyst 2960L-16PS-LL Switch
- Cisco Catalyst 2960L-16TS-LL Switch
- Cisco Catalyst 2960L-8PS-LL Switch
- Cisco Catalyst 2960L-8TS-LL Switch
- Cisco Catalyst 2960L-SM-48PQ Switch
- Cisco Catalyst 2960L-SM-48PS Switch
- Cisco Catalyst 2960L-SM-48TQ Switch
- Cisco Catalyst 2960L-SM-48TS Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 479 of 496

Cisco Catalyst 2960L-SM-24PQ Switch
Cisco Catalyst 2960L-SM-24PS Switch
Cisco Catalyst 2960L-SM-24TQ Switch
Cisco Catalyst 2960L-SM-24TS Switch
Cisco Catalyst 2960L-SM-16PS Switch
Cisco Catalyst 2960L-SM-16TS Switch
Cisco Catalyst 2960L-SM-8PS Switch
Cisco Catalyst 2960L-SM-8TS Switch

Cisco Catalyst 2960-Plus Series Switches

Cisco Catalyst 2960-Plus 48PST-L Switch
Cisco Catalyst 2960-Plus 48PST-S Switch
Cisco Catalyst 2960-Plus 48TC-L Switch
Cisco Catalyst 2960-Plus 48TC-S Switch
Cisco Catalyst 2960-Plus 24LC-L Switch
Cisco Catalyst 2960-Plus 24LC-S Switch
Cisco Catalyst 2960-Plus 24PC-L Switch
Cisco Catalyst 2960-Plus 24PC-S Switch
Cisco Catalyst 2960-Plus 24TC-L Switch
Cisco Catalyst 2960-Plus 24TC-S Switch

Cisco Catalyst 2960-S Series Switches

Cisco Catalyst 2960S-48FPD-L Switch
Cisco Catalyst 2960S-48FPS-L Switch
Cisco Catalyst 2960S-48LPD-L Switch
Cisco Catalyst 2960S-48LPS-L Switch
Cisco Catalyst 2960S-48TD-L Switch
Cisco Catalyst 2960S-48TS-L Switch
Cisco Catalyst 2960S-48TS-S Switch
Cisco Catalyst 2960S-24PD-L Switch
Cisco Catalyst 2960S-24PS-L Switch
Cisco Catalyst 2960S-24TD-L Switch
Cisco Catalyst 2960S-24TS-L Switch
Cisco Catalyst 2960S-24TS-S Switch

Cisco Catalyst 2960-SF Series Switches

Cisco Catalyst 2960S-F48FPS-L Switch
Cisco Catalyst 2960S-F48LPS-L Switch
Cisco Catalyst 2960S-F48TS-L Switch
Cisco Catalyst 2960S-F48TS-S Switch
Cisco Catalyst 2960S-F24PS-L Switch
Cisco Catalyst 2960S-F24TS-L Switch
Cisco Catalyst 2960S-F24TS-S Switch

Cisco Catalyst 2960-X Series Switches

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 480 of 496

Cisco Catalyst 2960X-48FPD-L Switch
Cisco Catalyst 2960X-48FPS-L Switch
Cisco Catalyst 2960X-48LPD-L Switch
Cisco Catalyst 2960X-48LPS-L Switch
Cisco Catalyst 2960X-48TD-L Switch
Cisco Catalyst 2960X-48TS-L Switch
Cisco Catalyst 2960X-48TS-LL Switch
Cisco Catalyst 2960X-24PD-L Switch
Cisco Catalyst 2960X-24PS-L Switch
Cisco Catalyst 2960X-24PSQ-L Cool Switch
Cisco Catalyst 2960X-24TD-L Switch
Cisco Catalyst 2960X-24TS-L Switch
Cisco Catalyst 2960X-24TS-LL Switch

Cisco Catalyst 2960-XR Series Switches

Cisco Catalyst 2960XR-48FPD-I Switch
Cisco Catalyst 2960XR-48FPS-I Switch
Cisco Catalyst 2960XR-48LPD-I Switch
Cisco Catalyst 2960XR-48LPS-I Switch
Cisco Catalyst 2960XR-48TD-I Switch
Cisco Catalyst 2960XR-48TS-I Switch
Cisco Catalyst 2960XR-24PD-I Switch
Cisco Catalyst 2960XR-24PS-I Switch
Cisco Catalyst 2960XR-24TD-I Switch
Cisco Catalyst 2960XR-24TS-I Switch

Cisco Cloud Services Platform 2100

Cisco Embedded Services 2020 Series Switches

Cisco Embedded Service 2020 24TC CON B Switch
Cisco Embedded Service 2020 24TC CON Switch
Cisco Embedded Service 2020 24TC NCP B Switch
Cisco Embedded Service 2020 24TC NCP Switch
Cisco Embedded Service 2020 CON B Switch
Cisco Embedded Service 2020 CON Switch
Cisco Embedded Service 2020 NCP B Switch
Cisco Embedded Service 2020 NCP Switch

Cisco Industrial Ethernet 3000 Series Switches

Cisco IE 3000-8TC Industrial Ethernet Switch
Cisco IE 3000-4TC Industrial Ethernet Switch

Cisco Industrial Ethernet 2000 Series Switches

Cisco IE 2000-24T67 Industrial Ethernet Switch
Cisco IE 2000-16T67 Industrial Ethernet Switch
Cisco IE 2000-8T67 Industrial Ethernet Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 481 of 496

Cisco ME 4600 Series Multiservice Optical Access Platform

Cisco ME 3800X Series Carrier Ethernet Switch Routers

Cisco ME 3800X-24FS-M Switch Router

Cisco ME 3600X Series Ethernet Access Switches

Cisco ME 3600X-24CX-M Switch
Cisco ME 3600X-24FS-M Switch
Cisco ME 3600X-24TS-M Switch

Cisco ME 3400E Series Ethernet Access Switches

Cisco ME 3400E-24TS-M Switch
Cisco ME 3400EG-12CS-M Switch
Cisco ME 3400EG-2CS-A Switch

Cisco ME 1200 Series Carrier Ethernet Access Devices

Cisco ME 1200-4S-A Ethernet Access Device
Cisco ME 1200-4S-D Ethernet Access Device

Cisco Meraki Cloud Managed Switches

Cisco Meraki MS420-48
Cisco Meraki MS420-24

Cisco Nexus 9000 Series Switches

Cisco Nexus 93180LC-EX Switch

Cisco Nexus 93180YC-EX Switch

Cisco Nexus 93180YC-EX-24 Switch
Cisco Nexus 93128TX Switch

Cisco Nexus 93108TC-EX Switch
Cisco Nexus 93108TC-EX-24 Switch
Cisco Nexus 9396PX Switch

Cisco Nexus 9396TX Switch

Cisco Nexus 9372PX Switch

Cisco Nexus 9372PX-E Switch
Cisco Nexus 9372TX Switch

Cisco Nexus 9372TX-E Switch
Cisco Nexus 9336PQ ACI Spine Switch

Cisco Nexus 9332PQ Switch

Cisco Nexus 7000 Series Switches

Cisco Nexus 7000 18-Slot Switch

Cisco Nexus 7000 10-Slot Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 482 of 496

Cisco Nexus 7000 9-Slot Switch

Cisco Nexus 7000 4-Slot Switch

Cisco Nexus 6000 Series Switches

Cisco Nexus 6004 Switch
Cisco Nexus 6001 Switch

Cisco Nexus 5000 Series Switches

Cisco Nexus 56128P Switch
Cisco Nexus 5696Q Switch
Cisco Nexus 5672UP Switch
Cisco Nexus 5672UP-16G Switch
Cisco Nexus 5648Q Switch
Cisco Nexus 5624Q Switch
Cisco Nexus 5596T Switch
Cisco Nexus 5596UP Switch
Cisco Nexus 5548UP Switch

Cisco Nexus 3000 Series Switches

Cisco Nexus 34180YC Switch
Cisco Nexus 31128PQ Switch

Cisco Nexus 3548 Switch

Cisco Nexus 3548-X Switch

Cisco Nexus 3524 Switch

Cisco Nexus 3524-X switch

Cisco Nexus 3464C Switch
Cisco Nexus 3264C-E Switch

Cisco Nexus 3264Q Switch

Cisco Nexus 3172PQ Switch

Cisco Nexus 3172PQ-XL Switch
Cisco Nexus 3172TQ Switch
Cisco Nexus 3172TQ-32T Switch
Cisco Nexus 3172TQ-XL Switch
Cisco Nexus 3164Q Switch

Cisco Nexus 3132C-Z Switch
Cisco Nexus 3132Q-XL Switch
Cisco Nexus 3064 Switch

Cisco Nexus 3064-T Switch
Cisco Nexus 3048 Switch

Cisco Nexus 2000 Series Fabric Extenders

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 483 of 496

Cisco Nexus 2348TQ 10GE Fabric Extender

Cisco Nexus 2348UPQ 10GE Fabric Extender

Cisco Nexus 2248PQ 10GE Fabric Extender

Cisco Nexus 2248TP GE Fabric Extender

Cisco Nexus 2232TM 10GE Fabric Extender

Cisco Nexus 2224TP GE Fabric Extender

Cisco Nexus 1100 Series Cloud Services Platforms

Cisco Nexus 1100 Cloud Services Platform

Cisco Nexus 1000V Switch for KVM

Cisco Nexus 1000V Switch for Microsoft Hyper-V

Cisco Nexus 1000V Switch for VMware vSphere

Cisco Nexus 1000V Switch

Cisco Nexus 1000VE

Cisco Small Business 500 Series Stackable Managed Switches

Cisco SF500-48 48-Port 10/100 Stackable Managed Switch
Cisco SF500-48MP 48-port 10/100 Max PoE+ Stackable Managed Switch
Cisco SF500-48P 48P-Port 10/100 POE Stackable Managed Switch
Cisco SF500-24 24-Port 10/100 Stackable Managed Switch
Cisco SF500-24MP 24-port 10/100 Max PoE+ Stackable Managed Switch
Cisco SF500-24P 24-Port 10/100 POE Stackable Managed Switch
Cisco SG500-52 52-port Gigabit Stackable Managed Switch
Cisco SG500-52MP 52-port Gigabit Max PoE+ Stackable Managed Switch
Cisco SG500-52P 52-port Gigabit POE Stackable Managed Switch
Cisco SG500-52PP 52-port Gigabit Max PoE+ Stackable Managed Switch
Cisco SG500-28 28-port Gigabit Stackable Managed Switch
Cisco SG500-28MPP 28-port Gigabit Max PoE+ Stackable Managed Switch
Cisco SG500-28P 28-port Gigabit POE Stackable Managed Switch
Cisco SG500-28PP 28-port Gigabit Max PoE+ Stackable Managed Switch
Cisco SG500X-48 48-Port GB with 4-Port 10-GB Stackable Managed Switch
Cisco SG500X-48MPP 48-port Gig Plus 4 10-Gig Max PoE+ Switch
Cisco SG500X-48P 48-P GB POE with 4-P 10-GB Stackable Managed Switch
Cisco SG500X-24 24-Port GB with 4-Port 10-GB Stackable Managed Switch
Cisco SG500X-24MPP 24-port Gig Plus 4 10-Gig Max PoE+ Switch
Cisco SG500X-24P 24P GB POE with 4Port 10GB Stackable Managed Switch
Cisco SG500XG-8F8T 16-port 10-Gigabit Stackable Managed Switch

Cisco Small Business 300 Series Managed Switches

Cisco SF302-08 8-Port 10/100 Managed Switch with Gigabit Uplinks
Cisco SF302-08MPP 8-port 10/100 Max PoE+ Managed Switch
Cisco SF302-08PP 8-port 10/100 PoE+ Managed Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 484 of 496

Cisco SF300-48 48-Port 10/100 Managed Switch with Gigabit Uplinks
Cisco SF300-48PP 48-port 10/100 PoE+ Managed Switch with Gig Uplinks
Cisco SF300-24 24-Port 10/100 Managed Switch with Gigabit Uplinks
Cisco SF300-24MP 24-port 10/100 Max-PoE Managed Switch
Cisco SF300-24PP 24-port 10/100 PoE+ Managed Switch with Gig Uplinks
Cisco SF300-08 8-Port 10/100 Managed Switch
Cisco SG300-52 52-Port Gigabit Managed Switch
Cisco SG300-52MP 52-port Gigabit Max-PoE Managed Switch
Cisco SG300-52P 52-port Gigabit PoE Managed Switch
Cisco SG300-28 28-Port Gigabit Managed Switch
Cisco SG300-28MP 28-port Gigabit Max-PoE Managed Switch
Cisco SG300-28PP 28-port Gigabit PoE+ Managed Switch
Cisco SG300-28SFP 28-port Gigabit SFP Managed Switch
Cisco SG300-20 20-Port Gigabit Managed Switch
Cisco SG300-10 10-Port Gigabit Managed Switch
Cisco SG300-10MPP 10-port Gigabit Max PoE+ Managed Switch
Cisco SG300-10PP 10-port Gigabit PoE+ Managed Switch
Cisco SG300-10SFP 10-port Gigabit Managed SFP Switch

Cisco Small Business 200 Series Smart Switches

Cisco SF200-48 48-Port 10/100 Smart Switch
Cisco SF200-48P 48-Port 10/100 PoE Smart Switch
Cisco SF200-24 24-Port 10/100 Smart Switch
Cisco SF200-24FP 24-port 10/100 Full-PoE Smart Switch
Cisco SF200-24P 24-Port 10/100 PoE Smart Switch
Cisco SG200-50 50-port Gigabit Smart Switch
Cisco SG200-50FP 50-port Gigabit Full-PoE Smart Switch
Cisco SG200-50P 50-port Gigabit PoE Smart Switch
Cisco SG200-26 26-port Gigabit Smart Switch
Cisco SG200-26FP 26-port Gigabit Full-PoE Smart Switch
Cisco SG200-26P 26-port Gigabit PoE Smart Switch
Cisco SG200-18 18-port Gigabit Smart Switch
Cisco SG200-10FP 10-Port PoE Smart Switch
Cisco SG200-08 8-Port Gigabit Smart Switch
Cisco SG200-08P 8-Port Gigabit POE Smart Switch

Cisco Small Business 110 Series Unmanaged Switches

Cisco SF112-24 24-Port 10/100 Switch with Gigabit Uplinks
Cisco SF110-24 24-Port 10/100 Switch
Cisco SF110-16 16-Port 10/100 Switch
Cisco SF110D-16 16-Port 10/100 Desktop Switch
Cisco SF110D-16HP 16-Port 10/100 PoE Desktop Switch
Cisco SF110D-08 8-Port 10/100 Desktop Switch
Cisco SF110D-08HP 8-Port 10/100 PoE Desktop Switch
Cisco SF110D-05 5-Port 10/100 Desktop Switch

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 485 of 496

Cisco SG112-24 Compact 24-Port Gigabit Switch
Cisco SG110-24 24-Port Gigabit Switch
Cisco SG110-24HP 24-Port PoE Gigabit Switch
Cisco SG110-16 16-Port Gigabit Switch
Cisco SG110-16HP 16-Port PoE Gigabit Switch
Cisco SG110D-08 8-Port Gigabit Desktop Switch
Cisco SG110D-08HP 8-Port PoE Gigabit Desktop Switch
Cisco SG110D-05 5-Port Gigabit Desktop Switch

Cisco Small Business 95 Series Unmanaged Switches

Cisco SF95-24 24-Port 10/100 Switch
Cisco SF95D-16 16-Port 10/100 Desktop Switch
Cisco SF95D-08 8-Port 10/100 Desktop Switch
Cisco SF95D-05 5-Port 10/100 Desktop Switch
Cisco SG95-24 Compact 24-Port Gigabit Switch
Cisco SG95-16 16-Port Gigabit Desktop Switch
Cisco SG95D-08 8-Port Gigabit Desktop Switch
Cisco SG95D-05 5-Port Gigabit Desktop Switch

Cisco Switch Modules for IBM

Cisco Catalyst Switch Module 3110 for IBM BladeCenter

Cisco Virtual Security Gateway

Cisco VSG for Microsoft Hyper-V
Cisco VSG for VMware vSphere

Citrix NetScaler 1000V

Rockwell Armor Stratix Series Switches

ARMORSTRATIX 5700 1783-ZMS24TA
ARMORSTRATIX 5700 1783-ZMS16TA
ArmorStratix 5700 1783-ZMS8TA

Unified Communications

Cisco ATA 190 Series Analog Telephone Adapters

Cisco ATA 190 Analog Telephone Adapter

Cisco Billing and Measurements Server

Cisco Business Edition 7000

Cisco Business Edition 7000 Version 12.0
Cisco Business Edition 7000 Version 11.6
Cisco Business Edition 7000 Version 11.5
Cisco Business Edition 7000 Version 10.6

Cisco Business Edition 6000

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 486 of 496

Cisco Business Edition 6000 Version 12.0
Cisco Business Edition 6000 Version 11.6
Cisco Business Edition 6000 Version 11.5
Cisco Business Edition 6000 Version 10.6
Cisco Business Edition 6000 Version 10.0

Cisco Business Edition 4000

Cisco Business Edition 4000 Appliance

Cisco Carrier Sensitive Route Server Software
Cisco Digital Gateway DE-30+
Cisco Digital Gateway DT-24+
Cisco Emergency Responder

Cisco Emergency Responder 12.0
Cisco Emergency Responder 10.5
Cisco Emergency Responder 10.0

Cisco Hosted Collaboration Solution for Contact Center

Cisco Hosted Collaboration Solution for Contact Center Version 11.5(1)
Cisco Hosted Collaboration Solution for Contact Center Version 9.0(1)

Cisco IAD2400 Series Integrated Access Devices

Cisco IAD2424 Integrated Access Device
Cisco IAD2423 Integrated Access Device
Cisco IAD2421-Integrated Access Device
Cisco IAD2420 Integrated Access Device

Cisco Prime Collaboration

Cisco Prime Collaboration Assurance 12.1
Cisco Prime Collaboration Assurance 11.6
Cisco Prime Collaboration Provisioning 12.6
Cisco Prime Collaboration Provisioning 12.5
Cisco Prime Collaboration Provisioning 12.4
Cisco Prime Collaboration Provisioning 12.3
Cisco Prime Collaboration Provisioning 12.2
Cisco Prime Collaboration Provisioning 12.1

Cisco Small Business Voice Gateways and ATAs

Cisco SPA8800 IP Telephony Gateway with 4 FXS and 4 FXO Ports

Cisco SPA8000 8-port IP Telephony Gateway

Cisco SPA122 ATA with Router

Cisco SPA112 2-Port Phone Adapter

Cisco WRP500 Wireless-AC Broadband Router with 2 Phone Ports

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 487 of 496

Cisco SVX4310 Software
Cisco TelePresence ISDN Gateway

Cisco TelePresence ISDN GW 3241
Cisco TelePresence ISDN GW MSE 8321

Cisco TelePresence Serial Gateway Series

Cisco TelePresence Serial GW 3340
Cisco TelePresence Serial GW MSE 8330

Cisco TelePresence Video Communication Server (VCS)

Cisco TelePresence Video Communication Server Control
Cisco TelePresence Video Communication Server Expressway
Cisco TelePresence Video Communication Server Model

Cisco Unified Communications for RTX
Cisco Unified Communications Manager (CallManager)

Cisco Unified Communications Manager Version 12.0

Cisco Unified Communications Manager Version 10.5

Cisco Unified Communications Manager Version 10.0

Cisco Unified Communications Manager Express

Cisco Unified Communications Manager Express Version 12.6
Cisco Unified Communications Manager Express Version 12.5
Cisco Unified Communications Manager Express Version 12.3
Cisco Unified Communications Manager Express Version 12.2
Cisco Unified Communications Manager Express Version 12.1
Cisco Unified Communications Manager Express Version 12.0
Cisco Unified Communications Manager Express Version 11.5

Cisco Unified Communications Manager IM & Presence Service

Cisco Unified Communications Manager IM and Presence Service 12.0
Cisco Unified Communications Manager IM and Presence Service Ver 10.5
Cisco Unified Communications Manager IM and Presence Service Ver 10.0

Cisco Unified PhoneProxy
Cisco Unified SIP Proxy

Cisco Unified SIP Proxy Version 10
Cisco Unified SIP Proxy Version 9.1
Cisco Unified SIP Proxy Version 9.0

Cisco Unity Express

Cisco Unity Express Version 10
Cisco Unity Express Version 9
Cisco Unity Express Version 8.6
Cisco Unity Express Version 8.5

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 488 of 496

Cisco uOne
Cisco VCO/4K Software
Cisco VG Series Gateways

- Cisco VG204XM Analog Voice Gateway
- Cisco VG202XM Analog Voice Gateway

Cisco Voice Network Switching System

Video

65/86 MHz Split

- GainMaker High Gain Balanced Triple Amp 1 GHz with 65/86 MHz Split
- GainMaker High Gain Dual System Amplifier 1 GHz with 65/86 MHz Split
- GainMaker Line Extender 1GHz with 65/86 MHz Split
- GainMaker Low Gain Dual System Amplifier 1 GHz with 65/86 MHz Split
- GainMaker Unbalanced Triple System Amp1 GHz with 65/86 MHz Split

55/70 MHz Split

- GainMaker High Gain Balanced Triple Amp 1 GHz with 55/70 MHz Split
- GainMaker High Gain Dual System Amplifier 1 GHz with 55/70 MHz Split
- GainMaker Line Extender 1GHz with 55/70 MHz Split
- GainMaker Low Gain Dual System Amplifier 1 GHz with 55/70 MHz Split
- GainMaker Unbalanced Triple System Amp 1 GHz with 55/70 MHz Split

42/54 MHz Split

- GainMaker High Gain Balanced Triple Amp 1GHz with 42/54 MHz Split
- GainMaker High Gain Dual System Amplifier 1GHz with 42/54 MHz Split
- GainMaker Line Extender 1GHz with 42/54 MHz Split
- GainMaker Low Gain Dual System Amplifier 1 GHz with 42/54 MHz Split
- GainMaker Unbalanced Triple System Amp 1 GHz with 42/54 MHz Split

40/52 MHz Split

- GainMaker High Gain Balanced Triple Amp 1GHz with 40/52 MHz Split
- GainMaker High Gain Dual System Amplifier 1GHz with 40/52 MHz Split
- GainMaker Line Extender 1GHz with 40/52 MHz Split
- GainMaker Low Gain Dual System Amplifier 1 GHz with 40/52 MHz Split
- GainMaker Unbalanced Triple System Amp 1 GHz with 40/52 MHz Split

Cisco Coaxial Media Converters

- Cisco 16x4 Coaxial Media Converters

Cisco Compact Amplifiers

- Cisco Compact EGC Dual Output Amplifier A93270
- Cisco Compact EGC Single Output Amplifier A93280
- Compact EGC Amplifier Model 93250
- Compact EGC Mini Amplifier A93230/A93240
- Reverse Amplifier Module A93146

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 489 of 496

Cisco Compact Nodes

Cisco Compact EGC Fiber Deep Node A90100/A90300
Cisco Compact GaN EGC Segmentable Node A90201
Compact Reverse Transmitters 9008x with FP, DFB or CWDM Lasers

Cisco Digital Media Players

Cisco Edge 340 Digital Media Player
Cisco Edge 300 Digital Media Player

Cisco GainMaker Amplifiers

Cisco 1 GHz GainMaker Amplifier

Cisco GainMaker Nodes

GainMaker Opto Node 1GHz with 40/52 MHz Split and RF Redundancy
GainMaker Optoelectronic Node 1GHz with 65/86 MHz Split
GainMaker Optoelectronic Node 1GHz with 55/70 MHz Split
GainMaker Optoelectronic Node 1GHz with 42/54 MHz Split
GainMaker Optoelectronic Node 1GHz with 40/52 MHz Split
GainMaker Reverse Segmentable Node - 1GHz with 40/52 MHz Split

Cisco GainStar Amplifiers

Cisco 862 MHz GainStar Amplifier
Cisco 1 GHz GainStar Amplifier

Cisco GainStar Nodes

Cisco 1 GHz GainStar Node

Cisco GS7000 Nodes

Cisco GS7000 4-Port 1GHZ Node
Model GS7000 4-Port Node 1 GHz with 65/86 Split
Model GS7000 4-Port Node 1 GHz with 55/70 Split
Model GS7000 4-Port Node 1 GHz with 42/54 Split

Cisco Hybrid Fiber Coax Configuration Tools

Cisco Handheld Programmer Terminal Model 91200

Cisco Line Equalizers

750 MHz In-line Equalizer
Cisco Multimedia Line Equalizer/Reverse Conditioner 1 GHz-65/86 Split
In-Line Equalizers
Multimedia Line Equalizer/Rev Conditioner 750 MHz - 42/51 MHz Split
Multimedia Line Equalizer/Reverse Conditioner 870 MHz - 42/51 Split
Multimedia Line Equalizer/Reverse Conditioner 1 GHz - 42/54 Split
Multimedia Line Equalizer/Reverse Conditioner 1 GHz - 40/52 Split
Multimedia Line Equalizer/Reverse Conditioner 1 GHz-85/105 Split

Cisco Prisma II Products

Prisma II Redundancy Interface Panel

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 490 of 496

Cisco RF Gateway Series

Cisco RF Gateway 10

Cisco RF Switches

Cisco Splitters, Directional Couplers, Power Inserters

Cisco Surge-Gap Passives

Cisco Stretch (Wide) Taps

Surge-Gap Stretch Taps

Cisco Traditional Size Tap Power Passing Accessories

Power Distribution Unit

Cisco Traditional Size Taps

1 GHz Surge-Gap Reverse Window Taps
1GHz Surge-Gap Taps, Standard & Full Profile
Flexible Solutions Taps

Cisco uBR10000 Series Universal Broadband Routers

Cisco uBR10012 Universal Broadband Router

Cisco uBR7200VXR Universal Broadband Routers

Cisco uBR7225VXR Universal Broadband Router

Cisco Vision

Cisco Vision Dynamic Signage Director

RF Accessories

Automatic Gain Control Modules
Interstage Equalizers
Interstage Trim Networks
Plug-in Pads, Forward and Reverse Equalizers
RF Amplifier Accessories
Thermal Level Control Modules

RF Signal Management

Series 9900 RF Signal Manager Active Products (NTSC)

Transponders

Cisco DOCSIS Transponder for GS7000 Node
Cisco DOCSIS/EuroDOCSIS Transponder for Compact Amplifiers and Nodes
Cisco DOCSIS/EuroDOCSIS Transponder for GainMaker Nodes

Wireless

Cisco 8500 Series Wireless Controllers

Cisco 8540 Wireless Controller
Cisco 8510 Wireless Controller

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 491 of 496

Cisco 5700 Series Wireless LAN Controllers

Cisco 5760 Wireless LAN Controller

Cisco 5500 Series Wireless Controllers

Cisco 5520 Wireless Controller
Cisco 5508 Wireless Controller

Cisco 4400 Series Wireless LAN Controllers
Cisco 3500 Series Wireless Controllers

Cisco 3504 Wireless Controller

Cisco 2500 Series Wireless Controllers

Cisco 2504 Wireless Controller

Cisco 2100 Series Wireless LAN Controllers
Cisco Access Policy Server
Cisco Aironet 4800 Access Points

Cisco Aironet 4800 Access Point

Cisco Aironet 3800 Series Access Points

Cisco Aironet 3800e Access Point
Cisco Aironet 3800i Access Point
Cisco Aironet 3800p Access Point

Cisco Aironet 3700 Series Access Points

Cisco Aironet 3700e Access Point
Cisco Aironet 3700i Access Point
Cisco Aironet 3700p Access Point

Cisco Aironet 3600 Series

Cisco Aironet 3600e Access Point
Cisco Aironet 3600i Access Point

Cisco Aironet 3500 Series

Cisco Aironet 3500e Access Point
Cisco Aironet 3500i Access Point
Cisco Aironet 3500p Access Point

Cisco Aironet 2800 Series Access Points

Cisco Aironet 2800e Access Point
Cisco Aironet 2800i Access Point

Cisco Aironet 2700 Series Access Points

Cisco Aironet 2700e Access Point
Cisco Aironet 2700i Access Point

Cisco Aironet 2600 Series

Cisco Aironet 2600e Access Point
Cisco Aironet 2600i Access Point

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 492 of 496

Cisco Aironet 1850 Series Access Points

Cisco Aironet 1850e Access Points
Cisco Aironet 1850i Access Points

Cisco Aironet 1840 Series Access Points

Cisco Aironet 1840i Access Point

Cisco Aironet 1830 Series Access Points

Cisco Aironet 1830i Access Point

Cisco Aironet 1815 Series Access Points

Cisco Aironet 1815i Access Point
Cisco Aironet 1815m Access Point
Cisco Aironet 1815t Access Point
Cisco Aironet 1815w Access Point

Cisco Aironet 1810 Series OfficeExtend Access Points

Cisco Aironet 1810 OfficeExtend Access Point

Cisco Aironet 1810w Series Access Points

Cisco Aironet 1810w Access Point

Cisco Aironet 1800 Access Points
Cisco Aironet 1800 Series

Cisco Aironet 1800i Access Point

Cisco Aironet 1700 Series Access Points

Cisco Aironet 1700i Access Points

Cisco Aironet 1600 Series

Cisco Aironet 1600e Access Point
Cisco Aironet 1600i Access Point

Cisco Aironet 1570 Series

Cisco Aironet 1572EAC Outdoor Access Point
Cisco Aironet 1572EC Outdoor Access Point
Cisco Aironet 1572IC Outdoor Access Point

Cisco Aironet 1560 Series

Cisco Aironet 1562D Outdoor Access Point
Cisco Aironet 1562E Outdoor Access Point
Cisco Aironet 1562I Outdoor Access Point

Cisco Aironet 1550 Series

Cisco Aironet 1552C Outdoor Access Point
Cisco Aironet 1552CU Outdoor Access Point
Cisco Aironet 1552E Outdoor Access Point
Cisco Aironet 1552EU Outdoor Access Point

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 493 of 496

Cisco Aironet 1552H Outdoor Access Point

Cisco Aironet 1552I Outdoor Access Point

Cisco Aironet 1552S Outdoor Access Point

Cisco Aironet 1552WU Outdoor Access Point

Cisco Aironet 1540 Series

Cisco Aironet 1542D Outdoor Access Point

Cisco Aironet 1542E Outdoor Access Point

Cisco Aironet 1542I Outdoor Access Point

Cisco Aironet 1530 Series

Cisco Aironet 1530e Outdoor Access Point

Cisco Aironet 1530i Outdoor Access Point

Cisco Aironet 1500 Series

Cisco Aironet 1000 Series

Cisco Aironet 1000 Series Lightweight Access Point

Cisco Aironet 350 Series Bridges

Cisco Aironet 340 Series

Cisco Aironet 340 Access Point-Captured Antenna

Cisco Aironet 340 Access Point-RP-TNC

Cisco Aironet 340 Access Points

Cisco Aironet Wireless LAN Client Adapters

Cisco Aironet 802.11a/b/g CardBus Wireless LAN Client Adapter (CB21AG)

Cisco Aironet 802.11a/b/g PCI Wireless LAN Client Adapter (PI21AG)

Cisco Aironet 340 Wireless PC Card Adapter

Cisco Aironet 340 Wireless PCI/LMC Adapter

Cisco Aironet 5 GHz 54 Mbps Wireless LAN Client Adapter (CB20A)

Cisco Catalyst 9117AX Series Access Points

Cisco Catalyst 9117AXI Access Point

Cisco Context-Aware Software

Cisco Flex 7500 Series Wireless Controllers

Cisco Flex 7510 Wireless Controller

Cisco Mobile Client

Cisco Policy Suite for BNG

Cisco Policy Suite for Wi-Fi

Cisco Secure Services Client

Cisco Small Business 500 Series Wireless Access Points

Cisco WAP561 Wireless-N Dual Radio Selectable Band Access Point

Cisco WAP551 Wireless-N Single Radio Selectable Band Access Point

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 494 of 496

Cisco Small Business 300 Series Wireless Access Points

Cisco WAP371 Wireless-AC/N Access Point with Single Point Setup

Cisco WAP351 Wireless-N Dual Radio Access Point with 5-Port Switch

Cisco WAP321 Wireless-N Access Point with Single Point Setup

Cisco Small Business 100 Series Wireless Access Points

Cisco WAP131 Wireless-N Dual Radio Access Point with PoE

Cisco WAP125 Wireless-AC Dual Band Desktop Access Point with PoE

Cisco WAP121 Wireless-N Access Point with Single Point Setup

Cisco Ultra-Reliable Wireless Backhaul

Cisco FM4200 Mobi

Cisco Universal Small Cell 9000 Series

Cisco Universal Small Cell 9330

Cisco Universal Small Cell 8000 Series

Cisco Universal Small Cell 8838
Cisco Universal Small Cell 8738
Cisco Universal Small Cell 8718
Cisco Universal Small Cell 8438
Cisco Universal Small Cell 8338
Cisco Universal Small Cell 8088
Cisco Universal Small Cell 8088v
Cisco Universal Small Cell 8050 Enterprise Management System

Cisco Universal Small Cell 7000 Series

Cisco Universal Small Cell 3000 Series

Cisco Universal Small Cell 3330

Cisco Virtual Wireless Controller

Orckit Exhibit 2004
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 495 of 496

UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

ORCKIT CORPORATION, Plaintiff, v. CISCO SYSTEMS, INC., Defendant.	Civil Action No. 2:22-cv-276-JRG-RSP **JURY TRIAL DEMANDED**

ORDER GRANTING DEFENDANT'S MOTION FOR A STAY PENDING *INTER PARTES* REVIEW PROCEEDINGS ON ALL FOUR ASSERTED PATENTS

Before the Court is Defendant's Motion for a Stay Pending *Inter Partes* Review Proceedings on All Four Asserted Patents. The Court, having considered same, finds that the Motion should be GRANTED. Accordingly, it is hereby:

ORDERED that this case is stayed in its entirety until further Order of this Court.