UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/285,028	05/22/2014	Ygdal Naouri	P62688	2072

96162 7590 12/27/2019

Law Office of R. Alan Burnett, PS
c/o CPA Global
900 Second Avenue South, Suite 600
Minneapolis, MN 55402

EXAMINER	
MIAN, OMER S	

ART UNIT	PAPER NUMBER
2461	

NOTIFICATION DATE	DELIVERY MODE
12/27/2019	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

alan@patentlylegal.com
docketing@cpaglobal.com

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Ex parte YGDAL NAOURI, ROBERT O. SHARP, KENNETH G. KEELS, and ERIC W. MULTANEN

Appeal 2018-007825
Application 14/285,028
Technology Center 2400

Before ROBERT E. NAPPI, JOHNNY A. KUMAR, and LINZY T. McCARTNEY, *Administrative Patent Judges*.

McCARTNEY, *Administrative Patent Judge.*

DECISION ON APPEAL

Appellant[1] seeks review under 35 U.S.C. § 134(a) of the Examiner's final rejection of claims 1–19 and 21–26. We have jurisdiction under 35 U.S.C. § 6(b).

We affirm in part.

[1] Appellant identifies the real party in interest as Intel Corporation. Appeal Brief 3, filed April 13, 2018 ("Appeal Br.").

BACKGROUND

This patent application concerns data center congestion management for non-Transmission Control Protocol traffic. *See*, *e.g.*, Specification ¶ 17, filed May 22, 2014 ("Spec."). Claims 1, 11, and 18 are independent. Claim 1 illustrates the claimed subject matter:

> 1. A method for implementing congestion management of non-TCP (Transmission Control Protocol) traffic in a network, comprising:
>
> determining a non-congested transit latency for a path between a source end-node and a destination end-node, wherein the path traverses at least one switch and the non-congested transit latency is an overall transit latency for the path;
>
> measuring transit latencies for transfer of non-TCP packets or Ethernet frames encapsulating non-TCP packets along the path, each transit latency comprising an overall transit latency for the path;
>
> determining, for transfer of each of at least a portion of the non-TCP packets or Ethernet frames, whether to mark the path as congested or not congested based on a difference between the transit latency measured for the non-TCP packet or Ethernet frame and the non-congested transit latency determined for the path; and
>
> managing a rate at which the non-TCP packets are transmitted from the source end-node to be forwarded via the path to the destination end-node based as a function of a rate at which the path is marked as congested.

Appeal Br. 44.

REJECTIONS[2]

Claims	35 U.S.C. §	References
1, 3–9, 11, 18, 21, 26	103	Cafiero,[3] Dahlin[4]
2, 12, 19	103	Cafiero, Dahlin, Attar[5]
10, 13–17, 22–25	103	Cafiero, Dahlin, Corlett[6]

DISCUSSION

We have reviewed the Examiner's rejections and Appellant's arguments, and with the exception discussed below, we disagree with Appellant that the Examiner erred. As consistent with the discussion below, we adopt the Examiner's reasoning, findings, and conclusions on pages 3–38 of the Final Office Action, pages 4–15 of the Advisory Action mailed February 5, 2018 ("Advisory Act."), and pages 2–34 of the Examiner's Answer. We address Appellant's arguments in turn.

Claim 1

"*Determining a Non-Congested Transit Latency*"

Claim 1 recites "determining a non-congested transit latency for a path between a source end-node and a destination end-node, wherein the path traverses at least one switch and the non-congested transit latency is an overall transit latency for the path." Appeal Br. 44. Appellant contends that

[2] In the Final Office Action, the Examiner rejected the pending claims under 35 U.S.C. § 112(b). Final Office Action 2, mailed November 16, 2017 ("Final Act."). The Examiner withdrew this rejection in the Examiner's Answer. Examiner's Answer 2, mailed May 29, 2018 ("Ans.").
[3] Cafiero et al. (US 2006/0098681 A1; May 11, 2006).
[4] Dahlin et al. (US 2004/0064577 A1; April 1, 2004).
[5] Attar et al. (US 2011/0211449 A1; September 1, 2011).
[6] Corlett (US 2009/0161569 A1; June 25, 2009).

one of ordinary skill in the art would "not need to combine Cafiero with Dahlin to obtain [the recited] non-congested overall latency" apparently because Dahlin "on its own, can determine the non-congested transit latency for a round-trip by determining the minRTT" or minimum observed round trip time. Reply Brief 8, filed July 29, 2018 ("Reply Brief") (emphases omitted).

We find this argument unpersuasive. Although one of ordinary skill in the art may not have needed to modify *Dahlin* to include the recited non-congested transit latency, the Examiner found that *Cafiero* does not teach aspects of this latency. *See, e.g.*, Final Act 5, 29. The Examiner found that Cafiero discloses determining an age limit for latency control (referred to as T_L) that corresponds to the recited "determining a non-congested transit latency." *See, e.g.*, Final Act. 4. But the Examiner found that Cafiero does not teach that this age limit "is an overall transit latency for the path" as required by claim 1. *See* Final Act. 5, 29. To remedy this deficiency, the Examiner turned to Dahlin. *See* Final Act. 5–6, 29–30. The Examiner found that Dahlin discloses a minimum round trip time (minRTT) that teaches a non-congested transit latency that is an overall transmit latency for a path. *See* Final Act. 5, 29–30; *see also* Dahlin ¶ 15 ("An estimate of uncongested round trip time may be based on a minimum round trip time for a data packet that has been detected (e.g., within a specific time period)."). Appellant acknowledges that Dahlin's minimum round trip time is a non-congested transit latency. *See* Appeal Br. 30 (conceding that Dahlin's "minRTT could be used as a non-congested transit latency"); Reply Br. 8 (admitting that Dahlin "can determine the non-congested transit latency for a round-trip by determining the minRTT").

The Examiner concluded that it would have been obvious to replace Cafiero's age limit for latency control with a minimum roundtrip time to arrive at the claimed invention because doing so would provide more accurate latency measurements, among other advantages. *See* Ans. 25 (explaining that using a roundtrip time "gives a more accurate and not just an inferred measurement by giving a more direct determination of latency in terms of time in both direction"). We thus find unpersuasive Appellant's contention that one of skill in the art would not have needed to combine Cafiero with Dahlin.

"Measuring Transit Latencies"

Claim 1 also recites "measuring transit latencies for transfer of non-TCP packets or Ethernet frames encapsulating non-TCP packets along the path, each transit latency comprising an overall transit latency for the path." Appeal Br. 44. Appellant argues that Dahlin cannot teach measuring transit latencies for transfer of *non-TCP* packets or Ethernet frames encapsulating *non-TCP packets* as recited in claim 1 because Dahlin uses *TCP* to measure roundtrip times. *See* Appeal Br. 27, 30–31. For this reason, Appellant also contends that "it would not be possible" to combine Dahlin's and Cafiero's inventions to "measur[e] transit latencies for transfer of non-TCP packets or Ethernet frames encapsulating non-TCP packets along the path, each transit latency comprising an overall transit latency for the path." Appeal Br. 32; *see also* Appeal Br. 34.

We find these arguments unpersuasive. The Examiner found Cafiero discloses packet ages that correspond to the recited "measuring transit latencies for transfer of non-TCP packets or Ethernet frames encapsulating non-TCP packets along the path" but does not disclose the each transit

latency "compris[es] an overall transit latency for the path" as required by claim 1. *See, e.g.*, Final Act. 4, 28–29. But the Examiner found that Dahlin discloses a roundtrip time (RTT) that teaches transit latencies that comprise an overall transit latency for the path and concluded that it would have been obvious to replace Cafiero's packet ages with a round trip time to arrive at the claimed invention. *See, e.g.*, Final Act. 5–6; Ans. 9–10.

Contrary to Appellant's arguments, the Examiner did not conclude that it would have been obvious to incorporate Dahlin's use of TCP into Cafiero. As discussed above, the Examiner relied on Dahlin only for the concept of roundtrip times. *See* Final Act. 32 (explaining that the Examiner used Dahlin "only . . . to teach . . . latency measurement[s] based on round trip times" and did not modify Cafiero by "the entire invention of" Dahlin); Ans. 28 (explaining that Dahlin "is *only relied upon to teach* round trip times" (emphasis modified)). In any event, "[t]he test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference . . . . Rather, the test is what the combined teachings of the references would have suggested to those of ordinary skill in the art." *In re Keller*, 642 F.2d 413, 425 (CCPA 1981). Thus, Appellant's arguments that one of ordinary skill in the art could not have combined Cafiero's and Dahlin's inventions have not shown that the Examiner erred. *Cf. In re Etter*, 756 F.2d 852, 859 (Fed.Cir.1985) (en banc) ("Etter's assertions that Azure cannot be incorporated in Ambrosio are basically irrelevant, the criterion being not whether the references could be physically combined but whether the claimed inventions are rendered obvious by the teachings of the prior art as a whole.").

The Examiner also found that the roundtrip times disclosed by Dahlin "are generic to all packet based communication protocols" and "implementable in any packet[-]base[d] protocol." Ans. 28; *see also* Final Act. 32 (finding that "[d]etermining round-trip times of packets traversing a path in a communication networks is a well-known technique in the art"). Appellant contends that the Examiner "cites no factual basis or evidence for" these findings and has not shown how one of ordinary skill in the art could modify Cafiero with Dahlin's teachings "to determine a round trip time without using TCP." Reply Br. 10; *see also* Reply Br. 20.

We find these arguments unpersuasive. Dahlin discloses that "[a] round trip time may refer to an elapsed time between the time that a data packet is sent and the time that the acknowledgement of receipt of the data packet is received." Dahlin ¶ 14. Thus, to determine these roundtrip times in Cafiero's system, the system would have to send and receive packets, determine when the packets were sent and received, and determine differences in time. Cafiero teaches sending and receiving packets, timestamping packets when certain events occur (for example, when the packet enters a buffer), and determining differences between a timestamp and the current time. *See, e.g.*, Cafiero ¶¶ 6 (explaining that the words "packet" and "frame" are used interchangeably in Cafiero), 73–75 (discussing sending packets and receiving frames), 142 (discussing timestamping packets), 143 (discussing comparing a timestamp with a current time). Given these disclosures, we see no error in the Examiner's finding that those of ordinary skill in the art would have been able to measure roundtrip times in Cafiero. We also note that another reference,

Corlett, teaches calculating roundtrip times. *See, e.g.*, Corlett ¶¶ 15, 45, 98–100, Abstract.

"Determining . . . Whether to Mark the Path"

Claim 1 also recites "determining, for transfer of each of at least a portion of the non-TCP packets or Ethernet frames, whether to mark the path as congested or not congested based on a difference between the transit latency measured for the non-TCP packet or Ethernet frame and the non-congested transit latency determined for the path." Appeal Br. 44. Appellant argues that the Examiner's combination of Cafiero and Dahlin does not teach or suggest this limitation because the combination would result in comparing Cafiero's *age value* with Dahlin's *minRTT measurement*, which Appellant contends are "completely different measurements." Appeal Br. 33.

We disagree. As discussed above, the Examiner concluded that it would have been obvious to replace Cafiero's age for a packet with a round trip time and age limit for latency control with a minimum round trip time. *See*, *e.g.*, Final Act. 29–30; Ans. 5–10. Thus, the Examiner's combination of Cafiero and Dahlin would result in comparing a *round trip time* with a *minimum round trip time*, not an age value with a minimum round trip time. We thus find this argument unpersuasive.

Motivation to Combine and Reasonable Expectation of Success

Appellant argues that the "the Examiner's motivation to combine has nothing to do with the claimed" invention because the context of the invention generally involves "use in a data center or the like" and the Examiner's motivation concerns "the use of a self-tuning background replication layer." Appeal Br. 34, 35. Appellant asserts that this use would "greatly limit the usefulness of the combination, and it could not be applied

to managing data center traffic congestion on a general level." Appeal Br. 35.

We disagree. As acknowledged by Appellant, claim 1 is "not limited exclusively to . . . use in a data center or the like," Appeal Br. 34, and Appellant has provided no persuasive evidence or reasoning to support the assertion that using a self-tuning background replication layer would "greatly limit the usefulness of the combination" and "could not be applied to managing data center traffic congestion on a general level," Appeal Br. 35. More important, Appellant ignores that the Examiner provided other reasons to combine Cafiero's and Dahlin's teachings in the claimed manner. The Examiner found that the proposed combination would provide more accurate latency measurements and better information about the status of network traffic. *See* Ans. 25 (explaining that using a round trip time "gives a more accurate and not just an inferred measurement by giving a more direct determination of latency in terms of time in both direction[s]"); Final Act. 30 (explaining that using round trip times "provides a well-known technique of measuring the round-trip latency as opposed to a latency towards the destination only to give a better overall status of the network traffic in a two-way communication service"). This reasoning provides an adequate motivation to combine the teachings of the cited art in the claimed manner.

Appellant also argues there would have been no motivation to combine the teachings of Cafiero and Dahlin or a reasonable expectation of success in doing so because the references "involve totally different approaches to managing network congestion" and "it would be impossible for the proposed combination to meet the limitations" of claim 1. Appeal Br. 34, 35; *see also* Appeal Br. 33, 36. According to Appellant, combining these

different approaches to managing network congestion would make Cafiero inoperable because Cafiero's "use of buffer fill levels and ECN would be replaced." Appeal Br. 35.

We again disagree. Appellant's arguments assume that the Examiner's combination involves every aspect of Cafiero's and Dahlin's inventions. The Examiner repeatedly made clear that it does not. *See* Final Act. 32 (explaining that the Examiner used Dahlin "only . . . to teach . . . latency measurement[s] based on round trip times" and did not modify Cafiero by "the entire invention of" Dahlin); Ans. 28 (explaining that Dahlin "is *only relied upon to teach* round trip times and at least one measurement of round trip time being 'non-congested transit latency'" (emphasis modified)). And as explained above, those of ordinary skill in the art would have been able to determine roundtrip times in Cafiero's system and thus would have a reasonable expectation of success in making the Examiner's proposed modifications. As for Cafiero's buffer fill levels and ECN, the Examiner's combination does not involve replacing these elements, and Appellant has provided no persuasive evidence or reasoning that using a roundtrip time or a minimum roundtrip time in Cafiero's invention would require these modifications.

Finally, Appellant contends that one of ordinary skill in the art would not have looked to Dahlin because the claimed invention manages congestion of non-TCP traffic and Dahlin uses TCP. Appeal Br. 35. We understand Appellant to argue that Dahlin is not analogous to the claimed invention. For a reference to be analogous to a claimed invention, the reference must be "from the same field of endeavor, regardless of the problem addressed" or "reasonably pertinent to the particular problem with

which the inventor is involved." *In re Bigio*, 381 F.3d 1320, 1325 (Fed. Cir. 2004). Here, Dahlin is at least reasonably pertinent to the particular problem with which the inventor is involved. The particular problem faced by the inventor included developing improved data center congestion management for non-TCP traffic. *See* Spec. ¶ 17. Although Dahlin addresses improving TCP congestion management, *see, e.g.*, Dahlin ¶ 9, one of ordinary skill in the art looking to improve the congestion control of non-TCP traffic would have considered how other transmission protocols addressed congestion control.

In any event, the problem faced by the inventor also included providing the benefits of TCP congestion control to non-TCP traffic. *See, e.g.*, Spec. ¶ 3 ("TCP traffic [benefits] from a built-in end-to-end congestion management method that has been enhanced by several techniques. Data Center Transport Control Protocol (DCTCP) is the most recent and the most efficient congestion avoidance variant . . . . Unfortunately, DCTCP is not relevant for non-TCP traffic . . . ."), 19 ("In accordance with one aspect, the techniques are implemented in a manner that mimics the DCTCP algorithm for non-TCP protocols, which has been proven to provide an efficient end-to-end congestion management method."). Dahlin addresses modifying a TCP congestion control protocol to, among other things, "be more sensitive to congestion than traditional protocols." Dahlin ¶ 9. Because Dahlin provides an improved TCP congestion control protocol, Dahlin would have been reasonably pertinent to the problem of providing the benefits of TCP congestion control to non-TCP traffic.

We have considered Appellant's remaining arguments and find them unpersuasive. For at least the above reasons, we find Appellant's arguments unpersuasive and therefore sustain the Examiner's rejection of claim 1.

Claim 2

Claim 2 recites "inputting into a DCTCP algorithm a congestion marking status of the path in place of using a congestion marking status conveyed via an ECN-Echo flag as an input to the DCTCP algorithm." Appeal Br. 44. Appellant argues that the Examiner's combination of Cafiero and Dahlin does not teach or suggest this limitation because Cafiero employs ECN for virtual lanes, not paths. *See* Appeal Br. 41. Appellant also argues that the Examiner did not address "replacing the ECN-Echo flag with anything, much less the congestion marking status of the path" as required by claim 2. Appeal Br. 42.

Appellant has not persuaded us that the Examiner erred. The Examiner found Cafiero uses ECN for paths because Cafiero discloses that a packet age can "indicate[] that *the path* towards the destination of the packet is subject to congestion" and that ECN is "considered a *core-to-end notification*, because it is *sent by core device 1420 and received by NIC card 1465 of end-node 1440*." Ans. 30 (quoting Cafiero ¶¶ 143, 137 (emphases modified)), 31 (showing an annotated version of Figure 14 of Cafiero). As for replacing the ECN-Echo flag, the Examiner found Cafiero teaches this aspect of claim 2 because Cafiero discloses setting the backward ECN portion (BECN) bit to indicate congestion. *See* Ans. 32 (quoting Cafiero ¶ 75). Appellant has not explicitly addressed these findings. We therefore find Appellant's arguments unpersuasive and thus sustain the Examiner's rejection of claim 2.

Claims 3 and 4

Claim 3 recites "wherein the non-congested transit latency is an overall transit latency that is measured along a one-way path from the source end-node to the destination end-node, and each of the transit latencies for the transfer of the non-TCP packets or the Ethernet frames is an overall transit latency that is measured along the same one-way path." Appeal Br. 45. Appellant argues the parts of Dahlin cited by the Examiner do not teach or suggest this limitation because they discuss measuring rounding trip times and make "no reference to measurement of a one-way path." Appeal Br. 39. Appellant contends that one of ordinary skill in the art "would not interpret a measurement of a round trip time that did not separately measure the one way path from the source node to a destination node to teach this limitation." Reply Br. 22; *see also* Appeal Br. 30.

We agree with Appellant. The Examiner found that Dahlin teaches this limitation because Dahlin discloses measuring a minimum round trip time. *See* Final Act. 9 (citing Dahlin ¶¶ 15, 118); Ans. 27 (citing Dahlin ¶¶ 118–119). The Examiner reasoned that "a roundtrip path includes a one-way path" thus "any latency value determined for the round trip would be for a one-way path" and regardless "any latency value measured along the roundtrip path has to be measured along at least the one-way path." Ans. 27 (emphases omitted). But as argued by Appellant, one of ordinary skill in the art would not interpret a measurement of roundtrip time as a latency that is "measured along a one-way path from the source end-node to the destination end-node" when the roundtrip time measurement does not include a separately measured one-way time. *See, e.g.*, Spec. ¶¶ 31–32 (discussing

one-way and roundtrip time measurements). We therefore do not sustain the Examiner's rejection of claim 3 and claim 4, which depends from claim 3.

Claim 11

Claim 11 recites the following:

> A server apparatus, configured to be implemented as an end-node in a data center, the server apparatus having a processor operatively coupled to memory and operatively coupled to a network interface configured to support network communication using an Ethernet protocol, the apparatus further configured, upon operation, to:
>
> facilitate operation of the server apparatus as a source end-node;
>
> . . . .
>
> retrieve path congestion marking indicia received from the destination end-node for transfer of at least a portion of the Ethernet frames encapsulating non-TCP packets, the path congestion marking indicia for a given Ethernet frame identifying whether or not the source destination path was marked as congested for transfer of the Ethernet frame based on a difference between a transit latency for the source-destination path measured for the Ethernet frame by the destination end-node, and a non-congested transit latency determined for the source-destination path, each of the transit latency measured for the Ethernet frame by the destination end-node and the non-congested transit latency comprising an overall transit latency for the source-destination path; and
>
> manage a rate at which the non-TCP packets are transmitted outbound from the transmit port to be forwarded via the source-destination path to the destination end-node based as a function of a rate at which the path is marked as congested.

Appeal Br. 47–48.

Appellant argues that the Examiner's rejection of claim 11 does not address these limitations. *See* Appeal Br. 36–37. Appellant also contends

that Cafiero does not teach or suggest these limitations because Cafiero performs its operations at "a switch, not a server that is operated as a source end-node." Appeal Br. 37. Finally, Appellant asserts that the Examiner erred in rejecting claim 11 for the same reasons that the Examiner erred in rejecting claim 1. *See* Appeal Br. 36.

Appellant has not persuaded us that the Examiner erred. The Examiner addressed these limitations in an Advisory Action mailed after the Final Office Action issued. *See, e.g.*, Advisory Action 10–14, mailed February 5, 2018. Appellant has not explicitly addressed the Examiner's findings in the Advisory Action for claim 11 and has therefore not persuaded us that the Examiner erred. *See* Manual of Patent Examiner Procedure § 1205.02 ("An appellant's brief must present arguments responsive to every ground of rejection stated by the examiner in the Office action from which the appeal has been taken (*as modified by any advisory action and/or pre-appeal brief conference decision*)." (emphasis added)). And for the reasons discussed above, we find Appellant's arguments about claim 1 (which Appellant asserts also apply to claim 11) unpersuasive. Accordingly, we sustain the Examiner's rejection of claim 11.

Remaining Claims

Appellant does not present separate, persuasive arguments for the other claims that depend from claim 1 (claims 5–10), the claims that dependent from independent claim 11 (claims 12–17), or independent claim 18 and its dependent claims (claims 19–26). We therefore sustain the Examiner's rejections of these claims.

CONCLUSION

Claims Rejected	35 U.S.C. §	References	Affirmed	Reversed
1, 3–9, 11, 18, 21, 26	103	Cafiero, Dahlin	1, 5–9, 11, 18, 21, 26	3, 4
2, 12, 19	103	Cafiero, Dahlin, Attar	2, 12, 19	
10, 13–17, 22–25	103	Cafiero, Dahlin, Corlett	10, 13–17, 22–25	
Overall Outcome			1, 2, 5–19, 21–26	3, 4

No period for taking any subsequent action in connection with this appeal may be extended under 37 C.F.R. § 1.136(a). *See* 37 C.F.R. § 1.136(a)(1)(iv).

AFFIRMED IN PART