

NOTE: This disposition is nonprecedential.

**United States Court of Appeals
for the Federal Circuit**

**MICROSOFT CORPORATION, INTERNATIONAL
BUSINESS MACHINES CORPORATION,**
Appellants

v.

PARALLEL NETWORKS LICENSING, LLC,
Appellee

2016-2515, 2016-2517, 2016-2518, 2016-2519, 2016-2642,
2016-2644, 2016-2645, 2016-2646

Appeals from the United States Patent and Trade-
mark Office, Patent Trial and Appeal Board in Nos.
IPR2015-00483, IPR2015-00484, IPR2015-00485,
IPR2015-00486, IPR2015-01729, IPR2015-01731,
IPR2015-01732, IPR2015-01734.

Decided: December 1, 2017

CONSTANTINE L. TRELA, JR., Sidley Austin LLP, Chi-
cago, IL, argued for appellants. Appellant Microsoft
Corporation also represented by RICHARD ALAN
CEDEROTH; MICHAEL D. HATCHER, Dallas, TX; JOSHUA
JOHN FOUGERE, DANIEL HAY, JOSEPH A. MICALLEF, Wash-
ington, DC.

JOHN M. DESMARAIS, Desmarais LLP, New York, NY, for appellant International Business Machines Corporation. Also represented by JON TODD HOHENTHNER, KEVIN KENT MCNISH, JEFFREY SCOTT SEDDON, II.

CHRISTOPHER THOR BOVENKAMP, McKool Smith, PC, Dallas, TX, argued for appellee. Also represented by DOUGLAS AARON CAWLEY; KEVIN P. HESS, DUSTIN MARK HOWELL, Austin, TX.

Before DYK, SCHALL, and TARANTO, *Circuit Judges*.

TARANTO, *Circuit Judge*.

Parallel Networks Licensing, LLC owns U.S. Patents 5,894,554 and 6,415,335 (the Parallel Patents), which describe and claim systems for managing the handling of requests for World Wide Web pages having dynamic (changing) content. '554 patent, col. 2, lines 15–19, lines 20–32; '335 patent, col. 2, lines 15–19, lines 20–32. Parallel sued Microsoft Corp. and International Business Machines (IBM) for infringement. Microsoft and IBM subsequently filed petitions with the Patent and Trademark Office seeking inter partes reviews of all claims of the Parallel Patents under 35 U.S.C. §§ 311–319. The Patent Trial and Appeal Board, acting as the delegate of the PTO's Director under 37 C.F.R. § 42.4(a), instituted reviews of all challenged claims on anticipation and obviousness grounds, and it consolidated the proceedings into two, one for each of the Parallel Patents, with Microsoft and IBM both involved as petitioners. After conducting the reviews, the Board concluded that the petitioners failed to demonstrate anticipation or obviousness.

Microsoft and IBM appeal the Board's decisions, arguing that the Board erred in construing the claim term

“request” and in rejecting the arguments for unpatentability. We affirm the Board’s claim construction, but we vacate in part the determination regarding anticipation, vacate the Board’s determination regarding obviousness, and remand for further proceedings.

I

A

The Parallel Patents claim priority to an application filed in 1996. By the mid-1990s, the number of Web page requests had grown large. *See* ’554 patent, col. 1, lines 39–41, col. 4, lines 33–39.¹ If many requests came to a single-processor server simultaneously, the server could be overwhelmed, and the efficiency of fulfilling requests could suffer. *Id.*, col. 4, lines 38–51. The Parallel Patents address the problem of processing numerous dynamic Web page generation requests using a partitioned architecture. *Id.* at col. 4, lines 51–53. The partitioned architecture includes a Web server, interceptor, dispatcher, and multiple page servers. *Id.*, fig. 4. The basic method recited in the challenged claims is as follows: (1) a Web page request is sent to a Web server; (2) the request is intercepted; (3) an appropriate page server to process the request is determined; and (4) the request is routed to the appropriate page server. *See id.*, col. 4, lines 54–60, col. 5, lines 38–43 & fig. 4.

The page server is determined based on dynamic system resource information, *e.g.*, the data sources each page server can access and the number of requests each page server is processing. *Id.*, col. 5, line 49 through col. 6, line 19. By routing a request to a page server, the Web server

¹ The ’335 patent issued from a divisional application that claims priority to the ’554 patent’s underlying application; the specifications are nearly identical. Thus, we omit duplicate citations to the ’335 patent.

is free to “concurrently process other Web client requests.” *Id.*, col. 6, lines 20–25. The partitioned architecture allows the page servers and the Web server “to simultaneously process different requests, thus increasing the efficiency of the Web site.” *Id.*, col. 6, lines 24–28.

In July 2012, the PTO issued *ex parte* reexamination certificates for both of the Parallel Patents.² All of the original claims in the ’554 patent were canceled and new claims 12–49 were added. J.A. 67. Similarly, all of the original claims in the ’335 patent were canceled and new claims 30–85 were added. J.A. 88. For present purposes, claim 12 of the ’554 patent is representative. It reads:

A computer-implemented method for managing a dynamic Web page generation request to a Web server, said computer-implemented method comprising the steps of:

routing said request from said Web server to a selected page server, said selected page server receiving said request and releasing said Web server to process other requests, wherein said routing step further includes the steps of intercepting said request at said Web server, routing said request from said Web server to a dispatcher, and dispatching, by said dispatcher, said request to said selected page server;

processing said request, said processing being performed by said selected page server while said Web server concurrently processes said other requests; and

² The ’335 patent’s *ex parte* reexamination certificate issued on July 17, 2012, and the ’554 patent’s *ex parte* reexamination certificate issued on July 24, 2012. Corresponding certificates of correction issued on those same days. J.A. 66–74, J.A. 87–97.

dynamically generating a Web page by said selected page server in response to said request, said Web page including data dynamically retrieved from one or more data sources; and

wherein dispatching includes:

examining said request to make a selection of which page server should process said request from among a plurality of page servers that can each generate said Web page requested by said request;

selecting one of said plurality of page servers to dynamically generate said Web page;

wherein said selection is based on examining dynamic information regarding a load associated with each of said plurality of page servers; and

sending said request to said selected page server based on said examination.

J.A. 69 (certificate of correction). The dispatcher may reside on the same machine as the Web server or on a different machine. '554 patent, col. 5, lines 8–28.

B

In December 2014, Microsoft filed four petitions for inter partes reviews, together challenging all claims of the Parallel Patents under 35 U.S.C. §§ 102 and 103.³ As to the '554 patent, the Board instituted reviews of claims 12–19, 32, 34, 46 and 48 in IPR2015-00483 (IPR-483) and

³ The Parallel Patents claim priority to a 1996 application and are governed by the versions of §§ 102 and 103 that were in effect before amendment by the Leahy-Smith America Invents Act, Pub. L. No. 112-29, § 3(n)(1), 125 Stat. 284, 293 (2011).

claims 12, 20–31, 33, 35–45, 47 and 49 in IPR2015-00484. The Board consolidated those reviews into IPR-483, which involves claims 12–49 of the ’554 patent. As to the ’335 patent, the Board instituted reviews of claims 30–40, 43–53, and 56–85 in IPR2015-00485 (IPR-485) and claims 32, 33, 35–42, 45, 46, 48–55, 65, 69, 80, and 84 in IPR2015-00486. The Board consolidated those reviews in IPR-485, which involves claims 30–85 of the ’335 patent. In August 2015, IBM filed four petitions for inter partes reviews, all substantially similar to the Microsoft petitions. The Board instituted reviews on the IBM petitions and joined IBM as a petitioner to IPR-483 and IPR-485.

In its August 2016 Final Written Decision for IPR-483, the Board concluded that Microsoft and IBM (hereafter collectively “Microsoft”) failed to demonstrate the unpatentability of claims 12–49 of the ’554 patent over the cited prior art. *Microsoft Corp. v. Parallel Networks Licensing, LLC*, IPR2015-00483, 2016 WL 8944632, at *11 (P.T.A.B. Aug. 11, 2016) (*IPR-483 Final Decision*). The Board reached the same conclusion in IPR-485 for claims 30–85 of the ’335 patent. *Microsoft Corp. v. Parallel Networks Licensing, LLC*, IPR2015-00485, 2016 WL 8999702, at *10 (P.T.A.B. Aug. 11, 2016) (*IPR-485 Final Decision*).⁴

Microsoft appeals the Board’s Final Written Decisions under 35 U.S.C. §§ 141(c) and 319. We have jurisdiction under 28 U.S.C. § 1295(a)(4)(A).

II

The main prior-art reference considered by the Board for both anticipation and obviousness is a 1995 publica-

⁴ The Board’s Final Written Decisions in IPR-483 and IPR-485 are substantially similar. We will refer only to the decision in IPR-483, but our conclusions apply to both.

tion by Daniel Andresen and others: *SWEB: Towards a Scalable World Wide Web Server on Multicomputers*, Dept. of Computer Sci. Tech. Report TRCS95-17, U.C. Santa Barbara (Sept. 1995) (*SWEB*). The publication discloses a “scalable Web server” implemented on a “cluster of workstations and parallel machines” for the efficient processing of relatively large numbers of simultaneous Web page requests. *Id.* at 1–2. The system includes a number of distributed Web servers that can either (a) themselves act as a page server when appropriate or (b), based on dynamic system load information, identify another Web server that can more efficiently get the client the desired information and take steps to have that server do so. *Id.* *SWEB* states that its preferred process for the latter is “URL redirection,” which refers to the Uniform Resource Locator tool for finding Web resources. *Id.* at 10. Figure 6 in *SWEB*, reproduced below, illustrates that process. *Id.* at 11.




Figure 6: The URL redirection.

The Web client (C) sends a Website host name, determined from a URL, to a domain name server (DNS), asking for an Internet Protocol (IP) address for a Web server (S?) hosting the site. *Id.* at 5. In the figure, the initial server to which the DNS directs the client is server S0. That server, like others in the system, contains a

scheduler (dispatcher) that communicates with all other Web servers to exchange system load information. *Id.* at 9. When the client sends request *r* to server *S0*, the scheduler determines whether *S0* or some other server should supply the client the desired information. *Id.*

If, at the time, server *S1* is preferable to *S0* as a source for the client to get the desired information, “URL Redirection” is used to bring about a connection between the client and *S1*. *Id.* at 10. “[W]hen client *C* sends a request [*r*] to server *S0*, *S0* returns a rewritten URL *r'* and a response code indicating the information is located at *r'*. *C* then follows *r'* to retrieve the resulting file *f*.” *Id.* The client computer can take the needed steps without a new action by, or even the knowledge of, the human user of client *C*: “Most Net browsers and clients automatically query the new location, so redirection is virtually transparent to the user.” *Id.*

This URL redirection process, while preferred by *SWEB*, is not the only process described in *SWEB* for changing the server that will actually supply the client-desired information. Before describing URL redirection, *SWEB* states that, where a server other than the initially contacted one would be better for serving the client at the time, “[t]he best situation would be to modify the UNIX sockets package” so that the initially contacted server sends the request directly to the desired Web server. *Id.* at 10. *SWEB* explains, however, that such modification “is difficult since it requires a substantial modification of the UNIX operating system kernel.” *Id.* It then explains the advantages of URL redirection and also the disadvantages, which it indicates are small. *Id.* at 10–11.

The primary advantages of URL redirection, *SWEB* states, are simplicity and universal compatibility. *Id.* at 10. “The primary disadvantage of URL redirection in practice is the added overhead of an additional connect/pass request/parse/respond cycle after the redirection

occurs.” *Id.* at 11. *SWEB* evaluates the magnitude of the disadvantage. It sets forth a mathematical model based on “a very short reply going back to the client browser, who then automatically issues another request to the new server address.” *Id.* at 14. It reports experimental results showing added overhead of URL redirection on the order of a few milliseconds in an overall process that took multiple seconds to complete. *Id.* at 21. *SWEB* characterizes this added overhead as “insignificant.” *Id.* at 23.

III

A

Microsoft first argues that the Board, when evaluating anticipation and obviousness, erroneously modified its otherwise-proper claim construction. The only claim term at issue on appeal is “request” in the required “dynamic Web page generation request.” The Board expressly adopted the construction of “request” that the parties agreed on during the pending district court litigation: “a message that asks for a Web page.” *IPR-483 Final Decision*, 2016 WL 8944632, at *4.

When conducting its analysis of the alleged anticipation by *SWEB*, the Board explained that the claims of the Parallel Patents “are directed to the routing and processing of a single request.” *Id.* at *5. That conclusion is not meaningfully in dispute here. And in any event, we see no error in it.

What Microsoft argues is that, in its anticipation analysis of whether *SWEB* teaches the claim-required steps involving a single request, the Board changed the agreed-to simple construction of “request.” Microsoft contends that the Board imposed requirements on what can be a “request” based on what Microsoft describes as (1) the particular route the request travels; (2) the particular URL or destination specified; and (3) who or what the request contacts along its route. But the Board did

none of those things. Instead, it determined that, under the agreed-on construction, each time the client computer sends a message asking for a Web page, it has made a new request. *IPR-483 Final Decision*, 2016 WL 8944632, at *6. The Board stated “that it is the act of asking—an action taken by the client computer—that defines a request.” *Id.* (emphasis omitted). Contrary to Microsoft’s contention, that conclusion does not turn on the route taken, destination, or intermediate contacts.

Microsoft’s core complaint seems to be that the Board should have focused on what the client is asking for, not each individual act of asking, and with that focus should have treated even multiple acts of asking as a single “request” as long as the content is the same. We have two difficulties with this position.

First, at best (for Microsoft’s purposes), there is ambiguity in the agreed-on construction’s term “message”—whether it refers to substantive communicative content or to a discrete act of utterance to convey communicative content. But *if* there is such ambiguity even in the present computer context, it is apparent on the face of the term. Yet Microsoft agreed to the term “message” as the defining construction in district court, leaving it to the fact finder to apply it. And in the present proceedings, Microsoft never presented to the Board a request for further claim construction.

Second, we agree with the Board as to the best construction in the context of this patent. This patent concerns computers sending out communications that other machines must process, with both the sending and the processing requiring expenditure of precious time and scarce resources. In this context, the more natural understanding of “request,” as of “message,” is one denoting each discrete act of communicating, not one that equates

discrete utterances when they convey the same content. *IPR-483 Final Decision*, 2016 WL 8944632, at *6–8.⁵

Accordingly, we reject Microsoft’s assertion that the Board committed a claim-construction error.

B

“[T]he dispositive question regarding anticipation is whether one skilled in the art would reasonably understand or infer from the prior art reference’s teaching that every claim element was disclosed in that single reference.” *Dayco Prods., Inc. v. Total Containment, Inc.*, 329 F.3d 1358, 1368 (Fed. Cir. 2003) (internal quotation marks, alterations, and citation omitted). Anticipation is a question of fact, subject to review for substantial evidence. *In re Rambus*, 753 F.3d 1253, 1256 (Fed. Cir. 2014).

Microsoft challenged the Parallel Patents as anticipated by two separate methods disclosed in *SWEB*: (1) URL redirection; and (2) modifying the UNIX sockets package. We affirm the Board’s finding that *SWEB*’s URL redirection method does not anticipate the claims of the Parallel Patents. But the Board failed to address the allegation of anticipation by the UNIX sockets modification method. We therefore vacate the Board’s rejection of anticipation in part and remand for the Board to consider that ground of anticipation in the first instance.

⁵ Because the Parallel Patents expired on April 23, 2016, before the Board issued its Final Written Decisions, the Board properly applied judicial claim-construction principles to identify the best construction, rather than seeking to identify the broadest reasonable construction. See *In re Rambus Inc.*, 694 F.3d 42, 46 (Fed. Cir. 2012) (inter partes reexamination); *Black & Decker, Inc. v. Positec USA, Inc.*, 646 F. App’x 1019, 1024 (Fed. Cir. 2016) (inter partes review).

1

Substantial evidence supports the Board’s finding that the Web client in the URL redirection disclosure in *SWEB* makes two separate requests as that term is used in the Parallel Patents. One is made when the initial request (r) is sent to the Web server selected by the DNS server. Another is made when the request (r’) is sent to the second Web server identified by the scheduler. *IPR-483 Final Decision*, 2016 WL 8944632, at *8.

Putting aside Microsoft’s claim-construction argument, which we have already rejected, Microsoft makes only one response. It accepts the Board’s focus on the redirection from one server to another. But it points out that the client computer’s human user makes only one communication—gives the client computer an instruction only once—in the *SWEB* URL redirection method in that circumstance. Microsoft’s Opening Br. 44–45; Microsoft’s Reply Br. 13. But that is irrelevant under the Parallel Patents. The patents refer interchangeably (for these purposes) to a request, Web client request, Web browser request, and URL request. And, more generally, the patents make clear that the requesting entity is the user’s computer, *i.e.*, the Web client, not the client’s human user. ’554 patent, col. 4, line 55 (“Web client 200 issues a URL request”), col. 6, lines 29–30 (“Web page is . . . transmitted back to requesting Web client 200”).

We find no error in the Board’s determination that *SWEB*’s URL redirection method fails to anticipate the challenged claims of the Parallel Patents.

2

Microsoft challenges the Board’s rejection of anticipation by *SWEB* on a separate ground. It observes, correctly, that the Board failed to address Microsoft’s argument for anticipation based on *SWEB*’s disclosure—as an alternative to URL redirection—of modifying the UNIX

sockets package of a server that has received a request so that the server forwards the request to a second server without, as in URL redirection, involving the client in initiating contact with the second server. Such a process is a form of what Microsoft's expert (and the Board) called "request forwarding" as a name for what the Parallel Patents require. See *IPR-483 Final Decision*, 2016 WL 8944632, at *9 (quoting J.A. 11449). We agree with Microsoft that the Board erred in not addressing this argument.

Parallel could not show that the Board actually did address this separate basis of alleged anticipation. It argues, however, that Microsoft waived this ground of anticipation. We cannot so find. Microsoft stated in its petitions that *SWEB* discloses two methods for routing requests: URL redirection and UNIX socket modification. J.A. 228. Microsoft also clearly argued UNIX socket modification as an alternative ground for anticipation of the Parallel Patents' routing limitation. J.A. 231. The Board did not find that those clear statements were insufficient to entitle Microsoft to a decision on the issue, and Parallel has not identified a regulatory or comparable requirement for preservation that Microsoft failed to meet. We see no basis for Parallel's argument for finding a complete waiver of the issue.

We vacate the rejection of anticipation in part and remand for the Board to consider the issue of anticipation by *SWEB*'s disclosure of modifying the UNIX sockets package. The Board may consider whether Microsoft committed a waiver and, if so, the scope of the waiver. If the Board finds no waiver, it should consider whether Microsoft has met the required elements of anticipation—"disclosure of all elements of [the] claimed invention arranged as in the claim." *Connell v. Sears, Roebuck & Co.*, 722 F.2d 1542, 1548 (Fed. Cir. 1983). And, if necessary, the Board should address Parallel's contention that this disclosure in *SWEB* is not enabled—applying our law

that generally accords a presumption of enablement to printed-publication and patent prior art. *In re Antor Media Corp.*, 689 F.3d 1282, 1289, 1292 (Fed. Cir. 2012) (printed publications); *Amgen Inc. v. Hoechst Marion Roussel, Inc.*, 314 F.3d 1313, 1355 (Fed. Cir. 2003) (patents).

C

Microsoft challenged the claims of the Parallel Patents as unpatentable for obviousness over *SWEB* (and other prior art) even if *SWEB* did not anticipate the “request forwarding” required by the claims. In rejecting the obviousness challenge, the Board relied on its determination that Microsoft and its expert had “not articulated a sufficient reason why a person of ordinary skill in the art would have modified *SWEB* to include request forwarding.” *IPR-483 Final Decision*, 2016 WL 8944632, at *9–10. Microsoft argues on appeal that the Board failed to give proper effect to the testimony of its expert under the legal standards for obviousness. We agree to this extent: the Board provided insufficient explanation to justify rejecting the challenge. We vacate the rejection of the obviousness challenge and remand for further proceedings on that challenge. *See, e.g., In re Nuvasive, Inc.*, 842 F.3d 1376, 1383 (Fed. Cir. 2016); *Synopsys, Inc. v. Mentor Graphics Corp.*, 814 F.3d 1309, 1322 (Fed. Cir. 2016), *overruled on other grounds by Aqua Prod., Inc. v. Matal*, 872 F.3d 1290 (Fed. Cir. 2017); *In re Sang Su Lee*, 277 F.3d 1338, 1346 (Fed. Cir. 2002).

An obviousness challenger “must demonstrate . . . that a skilled artisan would have had reason to combine [or modify] the teaching of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success from doing so.” *Redline Detection, LLC v. Star Envirotech, Inc.*, 811 F.3d 435, 449 (Fed. Cir. 2015) (quoting *PAR Pharm., Inc. v. TWI Pharm., Inc.*, 773 F.3d 1186, 1193

(Fed. Cir. 2014)). A reason “to modify a reference can come from the knowledge of those skilled in the art, from the prior art reference itself, or from the nature of the problem to be solved.” *SIBIA Neurosciences, Inc. v. Cadus Pharm. Corp.*, 225 F.3d 1349, 1356 (Fed. Cir. 2000).

Those standards can and must be understood to follow the Supreme Court’s various formulations in *KSR Int’l Co. v. Teleflex, Inc.*, 550 U.S. 398 (2007). The *KSR* formulations include those referring to the “combination of familiar elements according to known methods” that “does no more than yield predictable results,” *id.* at 416; “a structure already known in the prior art that is altered by the mere substitution of one element for another known in the field,” requiring more than “a predictable result” for patentability, *id.*; “simply arrang[ing] old elements with each performing the same function it had been known to perform,’ and yield[ing] no more than one would expect from such an arrangement,” *id.* at 417; “a technique [that] has been used to improve one device,” where “a person of ordinary skill in the art would recognize that it would improve similar devices in the same way” and that application is not beyond the ordinary skilled artisan’s skill, *id.*; “the simple substitution of one known element for another or the mere application of a known technique to a piece of prior art ready for the improvement,” *id.*; circumstances in which “there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions” and “the known options [are] within [an ordinary skilled artisan’s] technical grasp” and pursuing those options “leads to the anticipated success,” in which case the options are “obvious to try” and “might” be obvious, *id.* at 421. That enumeration is not exhaustive, and we are not identifying what aspects of *KSR* are ultimately important in this case.

Microsoft’s expert, Dr. Mitzenmacher, stated in his initial declaration:

214. To the extent Patent Owner may argue that SWEB 95 does not disclose “*routing said request from said Web server to a selected page server*” considering either disclosed method of sending requests to other servers, it would have been obvious to include that functionality in the scheme of SWEB 95. It was known in the prior art to route requests from one machine to another via request forwarding based upon a determination of which server should be used to process such a request. *See, e.g.*, Berson, *Client/Server Architecture* (2d. ed. 1996) (published by April 11, 1996 per U.S. Copyright Office) (“Berson”) (Ex. 1036) at 207-208 [J.A. 13311–12] (“Client/server interactions are defined as those in which a client initiates the interaction and requests a particular service from its ‘home’ server. This server reacts to the client’s request by either performing the desired service or routing the request to another server for fulfillment of the request; irrespective of the final location of the service, the original or home server sends the response back to the client.”); Pfister, *In Search of Clusters* (1995) (“Pfister”) (Ex. 1037) at 78 (processing node in a cluster may “dynamically move work from one internal node to another to better balance the node”), 226-228 [J.A. 13926–28]; Ex. 1012 [Garland *et al.*, *Implementing Distributed Server Groups for the World Wide Web* (1995)] at 8 [J.A. 12427] (“We could also extend the dispatcher to support anonymous members. It is conceivable that some servers might not want to advertise their location to the rest of the world. In this case, the dispatcher could *serve as an intermediary between the client and the anonymous server.*”) (emphasis added). And performing such routing was known to free up CPU cycles and memory resources of the sending device – that is

why the requests are routed away. *See, e.g.*, Ex. 1036 [Berson] at 9; Ex. 1037 [Pfister] at 226-228 [J.A. 13926–28]; Ex. 1012 [Garland] at 1 [J.A. 12420] (“When the volume of requests to a server becomes prohibitively high the server tends to drop requests and eventually crashes. The service provider needs more computing power to meet the demands of the clients requesting the service. ... A *fairly obvious solution to this problem* is to implement a distributed Web server which will spread incoming request load among several machines.”) (emphasis added). A person of ordinary skill in the art would therefore have been motivated to employ such a technique in the scheme of SWEB 95, which is expressly concerned with the efficient use of server resources. *See* Ex. 1009 [SWEB] at 2 (“Several resource constraints affect the performance of the server: the CPU speed and memory size of one processing unit, the existing system load, the transmission bandwidth between the processing unit and its local disk, the network latency and bandwidth between a processing unit and a remote disk when the accessed files are not stored in the local disk, and disk contention when multiple I/O requests access the same disk, can all affect the performance. Server scalability is achieved by actively monitoring the run time CPU, disk IO, and network loads of system resource units, and dynamically scheduling user HTTP requests to a proper node for efficient processing.”). Moreover, to do so would have been within the level of ordinary skill in the art, as the prior art shows, and merely the application of a prior art technique for its known purpose to achieve a predictable result.

J.A. 11449–51.

The Board, referring to this testimony, concluded that Microsoft had not explained why a person of ordinary skill in the art would have been motivated to modify *SWEB*'s URL redirection, because, it determined, the stated benefits of request forwarding were already provided by URL redirection. *IPR-483 Final Decision*, 2016 WL 8944632, at *9. This reasoning is inadequate. It does not address the legal significance, in light of some of the *KSR* formulations, of Dr. Mitzenmacher's statement, which is amply supported by references of record, that request forwarding was a well-known technique for a request-receiving server to use to get another server to provide the client the desired information. It does not address the statement that an interest in anonymity is a reason for request forwarding. It does not address the statements in *SWEB* itself suggesting that the "best situation" might be to use a form of request forwarding (by UNIX socket modification) and identifying a "disadvantage" of URL redirection. *SWEB* at 10, 11.

In these circumstances, we conclude that the Board failed to provide adequate reasoning to support its determination on the component of the obviousness analysis it said was decisive. That inadequacy, together with the failure to consider the contention that *SWEB* teaches UNIX sockets modification, requires vacatur of all of the Board's obviousness rulings. We remand for reconsideration of obviousness.

IV

For the foregoing reasons, we affirm the Board's claim construction and rejection of anticipation by the URL redirection embodiment of *SWEB*. We vacate the rejection of anticipation based on the UNIX socket modification teaching in *SWEB* and the rejection of the obviousness challenge. We remand for further proceedings consistent with this opinion.

No costs.

**AFFIRMED IN PART, VACATED IN PART, AND
REMANDED**