UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE INC.
Petitioner

v.

AT HOME BONDHOLDERS' LIQUIDATING TRUST
Patent Owner

Case IPR No. Unassigned
U.S. Patent 6,286,045

Petition for *Inter Partes* Review of U.S. Patent No. 6,286,045 UNDER 35 U.S.C. §§ 311-319 and 37 C.F.R. §§ 42.1-.80, 42.100-.123

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

I. INTRODUCTION

Google Inc. ("**Petitioner**") petitions for *Inter Partes* Review, seeking cancellation of claims 1-12, 14-19, and 34-42 ("**challenged claims**") of U.S. Patent No. 6,286,045 to Griffiths *et al.* ("**the '045 patent**") (GOOG 1001), which is owned by At Home Bondholders' Liquidating Trust ("**Patent Owner**").

II. GROUNDS FOR STANDING (37 C.F.R. § 42.104(a)); PROCEDURAL STATEMENTS

Petitioner certifies that the '045 patent is available for IPR. Petitioner further certifies that it is not barred or estopped from requesting IPR of claims 1-12, 14-19, and 34-42 on the grounds identified in this petition, as Petitioner was first served less than a year ago with a complaint for infringement on Feb. 20, 2014, in U.S. District Court for the District of Delaware (1:14-cv-00216)[1]. (GOOG 1015.) Concurrently filed herewith are Powers of Attorney and an Exhibit List per § 42.10(b) and § 42.63(e), respectively. The required fee is paid via online credit card payment. The Office is authorized to charge fee deficiencies and credit overpayments to Deposit Acct. No. 19-0036 (Customer ID No. 45324).

---

[1] Petitioner notes that it was also served with a complaint based on the '045 patent on Feb. 10, 2014. However, that complaint was dismissed without prejudice, and is therefore not relevant to the IPR bar date. (IPR2012-00004, Paper No. 18.)

III. STATEMENT OF THE PRECISE RELIEF REQUESTED AND THE REASONS THEREFOR (37 C.F.R. § 42.22(a))

Petitioner requests IPR and cancellation of claims 1-12, 14-19, and 34-42 of the '045 patent. A detailed statement of the reasons for the relief requested is set forth in §§ IV and VII below.

IV. OVERVIEW

Inter partes review ("**IPR**") was created to improve patent quality and, if warranted, cancel unpatentable claims. That core purpose is furthered by this Petition, as the challenged claims of the '045 patent should never have been issued. Not only was the alleged invention known before the '045 patent filing date, but the four "fundamental principles" of the alleged invention – touted by the Patent Owner during prosecution as distinguishing the invention from the art – were also well-understood by the industry. Because Petitioner is, at a minimum, reasonably likely to prevail in showing unpatentability, the Petition should be granted and trial instituted on all of the challenged claims as set forth below.

A. The '045 Patent

The '045 patent was filed on May 19, 1997, and issued on September 4, 2001. According to USPTO assignment recordation records, At Home Bondholders' Liquidating Trust is now the Patent Owner.

The '045 patent claims nothing more than a well-known method of Internet advertising and the ability to accurately account for the number of times an

advertisement is displayed to a user, while reducing heavy Internet traffic. In general, the '045 patent's claims describe distributing a banner advertisement over a network. This involves generating a request for an ad banner from a user's computer. Rather than sending the request directly to an advertiser's web site, the request is sent to another server that counts the request and redirects the request to the desired advertiser's web site. The redirected address is sent to the user's computer and a request from the user's computer is then sent to the selected advertising web site.

B. Prosecution History

In arguments made during prosecution, Patent Owner emphasized four "fundamental principles" of the alleged invention that the Examiner was to keep in mind when analyzing the prior art. (GOOG 1002, p. 149.) First, "Applicants' invention reduces the inaccurate display counting caused by caching of the banners by making or causing request signals generated or transmitted by a client device unblockable by the client device or proxy server, even though the banners may have been previously stored on the client device or proxy server." (*Id.*, p. 150.) "Second, applicants' invention allows such serving and counting to occur without significantly increasing data traffic on the computer network or unnecessarily delaying the display of the banners or other information on the client device." (*Id.*, p. 151.) "Third, applicants' invention allows banners or advertisements to be

targeted to users to increase the banners' or advertisements' effectiveness." (*Id.*) "Fourth, applicants' invention increases fault tolerance and reliability for information and banner delivery and storage systems." (*Id.*) Such arguments were apparently persuasive to the Examiner in overcoming the cited prior art. Yet each of these "fundamental principles," along with the mechanisms described in the '045 patent as embodying those principles, were well known in the industry before the '045 patent was filed.

C. State of the Art

1. Serving and Counting of Banners was Well-Known

Paul Leach was an early member of the HTTP Working Group of the World Wide Web Consortium and, in the mid-1990s, was heavily involved in developing the protocols by which Web traffic was governed.[2] (GOOG 1005, ¶ 8.) In his declaration, Mr. Leach explains that "[t]he concept and concern for accurately counting the number of times a banner was displayed on a client device was a well known issue at the time of the filing of the '045 patent." (GOOG 1005, ¶ 20.) Mr. Leach has also explained that "[i]t was also well known that the use of cache would cause an underreporting of the counting of banners. ... 'A request is a connection to

---

[2] Mr. Leach was a contributor to both the HTTP 1.0 and HTTP 1.1 specifications. (GOOG 1008, pp. 41-42; GOOG 1026, pp. 99-100.)

an Internet site (*i.e.*, hit) that successfully retrieves content,' but counting such requests accurately was a known issue 'because browser software and many Internet gateways intercept some requests before reaching the server, and these cached requests are never logged.'" (GOOG 1005, ¶ 21 (quoting GOOG 1022, p. 13).)

Not only was the problem a known issue, but Patent Owner's solution was also known. Specifically, cache avoidance to reliably determine page views – also known as "cache-busting" – was well known by early 1997. As Mr. Leach stated in his co-authored HTTP Working Group paper, "[f]or a variety of reasons, content providers want to be able to collect information on the frequency with which their content is accessed. This desire leads to some of the 'cache-busting' done by existing servers. ('Cache-busting' is the use by servers of techniques intended to prevent caching of responses...)" (GOOG 1024, pp. 2-3.) Further, as discussed in an earlier version of the same Working Group paper, "[s]ome cache-busting is also done to provide different advertising images to appear on the same page (*i.e.*, each retrieval of the page sees a different ad)....HTTP/1.1 already allows origin servers to prevent caching of responses, and we have evidence that at least some of the time, **this is being done for the sole purpose of collecting counts of the number**

of accesses of specific pages."[3] (GOOG 1016, pp. 2-3.)

Peter Kent is another expert in the field, having been involved in Internet advertising from the early days of the Web. (GOOG 1003, ¶¶ 5-8.) Mr. Kent agrees that:

> [c]ounting accuracy for delivered content was a widely known issue at the time the '045 patent was filed, and the proposed solution in the '045 patent was also already widely known. In fact, attempting to improve counting accuracy via cache avoidance was such a burden on the Web's bandwidth that by early 1997, other proposals were already being made to move advertisers away from the use of cache-avoidance. At any rate, such cache avoidance was already widely known before the alleged invention. (GOOG 1003, ¶15.)

Further, the cache avoidance methods described in the '045 patent specification ("HTTP no-cache pragma, appending a random segment to the URL, and using a cgi script to generate dynamic pages") were already known before the patent's filing date. (GOOG 1003, ¶16.) Mr. Kent also describes additional known cache avoidance methods, including modifying URLs and the use of third party products such as PageMeter. (GOOG 1003, ¶17.)

Thus, it was well-known by a person of ordinary skill in the art ("**POSA**") when the '045 patent was filed "that caching distorted the accurate counting of the

---

[3] Emphasis added throughout unless otherwise noted.

display of advertising banners and web pages and that there were known 'cache-busting' methods that could be used to prevent caching and thus allow for a more accurate method of counting the delivery and display of Internet based advertisements." (GOOG 1005, ¶ 27.)

2. Serving and Counting Banners without Significantly Increasing Network Traffic was Well-Known

While cache-busting allowed for a more accurate counting of banners, it was recognized in the industry at the time that cache-busting "also resulted in increased loads on servers." (GOOG 1005, ¶ 28.) "Not only was this method expensive computationally to the sever, but it defeated intermediary caching and did not correctly handle the exchanging of URLs between people." (GOOG 1005, ¶ 28, GOOG 1023, p. 2.) So, methodologies to allow for the serving and counting of banner advertisements without significantly increasing data traffic through the use of cache-busting were also "well known at the time of the filing of the '045 patent application." (GOOG 1005, ¶29.)

Mr. Leach declares that "[o]ne of ordinary skill in the art at the time of the filing of the '045 patent would have known that cache control mechanisms such as If-Modified-Since or If-NoneMatch headers would result in not blocking the request signal from reaching the intended server, but avoiding a refetch of the requested information if that information existed in cache." (GOOG 1005, ¶32, *see*

also GOOG 1005, ¶¶33-34.)

Mr. Leach's declaration also discusses a "hit-metering" approach that he and Jeffrey Mogul developed. That approach "outlines a method of counting requests, or 'hit counts' without defeating the use of cache where appropriate." (GOOG 1005, ¶ 35.) As Mr. Leach puts it, "[o]ur hit-metering approach allowed content providers to be able to collect information on the frequency with which their content is accessed, but without resorting to 'cache-busting' techniques discussed above that defeat the use of cache." (*Id.*; *see also,* GOOG 1024.)

Thus, "it was well known by a POSA at the time of the filing of the '045 patent that multiple methods existed that would allow for the accurate counting of banner advertisement requests without significantly increasing data traffic and that also allows for the efficient use of cache." (GOOG 1005, ¶36.)

3. Advertisement Targeting based on Demographics was Well-Known

Mr. Leach explains that "[t]he concept of targeting advertisements to particular users to increase advertising effectiveness was a well known issue at the time of the filing of the '045 patent." (GOOG 1005, ¶37.) As discussed in Mr. Leach's Hit-Metering paper, "some advertisers employed the use of 'cache-busting' to 'collect demographic information' so that advertising images could be tailored and targeted to those demographics, *e.g.*, 'each retrieval of the page sees a different

ad.'" (GOOG 1005, ¶ 37 (quoting GOOG 1024, p. 3).)

And Mary Meeker, in her detailed analysis of Internet advertising in early 1997, noted, "the Internet offers the ability to target and deliver messages to an audience with specific demographics and interests." (GOOG 1010, p. 3-13.)

> "Targeting gives advertisers the opportunity to filter messages to selected audiences based on certain criteria. This may be the most powerful aspect of the Internet as an advertising medium — the ability to dictate the exact composition of an advertisement's audience…each individual delivery can be tailored, based on user information. The power of the second aspect is increased substantially with more detailed user data, potentially collected through registration or in the course of using the site." (GOOG 1010, p. 6-3.)

Thus, "it was well-known by a POSA at the time of the filing of the '045 patent that advertisers were developing methods using demographics to increase advertising effectiveness." (GOOG 1005, ¶ 39.)

4. Fault Tolerance and Reliability Were Well Known

Based on Mr. Leach's extensive experience in "ACM conferences and committees directed to distributed computing, replication and fault tolerance as early as 1985" and his "published papers on the theories and principles of distributed computing in 1982, 1985 and 1987," he explains that the "concept of fault tolerant computing for increased reliability was a well-known concept at the time of the filing of the '045 patent."(GOOG 1005, ¶ 40.). Indeed, "[m]irroring

and redundancy were common fault tolerant methods at the time." (*Id.*) Thus, a POSA at the time of the '045 patent would have understood "that fault tolerant solutions in distributed computing existed and provided increased reliability in computer delivery and storage systems." (GOOG 1005, ¶ 41.)

5. HTTP Redirect was Well-Known

Mr. Kent explains that "HTTP redirect was also widely employed in the field of information delivery (including delivery of online advertisements)," and points to prior art including "Wexler [that] describes a third party accounting and statistical service 'configured to issue a '302' redirect response when a specific URL is requested.'" (GOOG 1003, ¶ 18 (quoting GOOG 1007, 5:16-17).) Mr. Kent further notes that "Merriman describes an advertisement server for 'send[ing] the redirect message causing the user's browser to receive the URL for the advertiser's web site based upon data stored in the server.'" (GOOG 1003, ¶ 18 (quoting GOOG 1013, 7:22-26).)

Indeed, Mr. Kent explains that "[o]ne well-known use of HTTP redirect messages was to refer a client computer to a server located in the close geographical proximity of the client for reducing latency....HTTP redirect messages [were] to refer the client computer to a selected server in a group of distributed servers," because "a group of web servers can reduce latency because a distributed web server group can balance the load and dispatch the request to the

least loaded web server." (GOOG 1003, ¶ 19.)

A POSA, considering the '045 patent's claims in light of the prior art, would have understood that the prior art rendered the claims unpatentable. The prior art references are discussed in detail below in Sec. VII. Each applied reference is analogous art to the claimed invention at least because it (1) falls within the '045 patent's stated field of "storage, management, and delivery of information on a computer network" (GOOG 1001, 1:9-11), and/or (2) is reasonably pertinent to one of the apparent problems allegedly solved.

As such, the challenged claims are well known and should not have been issued. Instead they should be cancelled. In view of the showings of obviousness provided below, Petitioner is reasonably likely to prevail in establishing that each of claims 1-12, 14-19, 34-38, and 40-42 of the '045 patent is unpatentable.

V. CLAIM CONSTRUCTION

In accordance with 37 C.F.R. § 42.100(b), the challenged claims must be given their broadest reasonable interpretations ("**BRI**") in light of the specification of the '045 patent. The following terms and phrases from the claims of the '045 patent require construction in accordance with these principles for the purpose of this IPR. The plain and ordinary meaning should be applied to any claim terms that are not addressed below. Petitioner reserves the right to pursue different constructions under different standards applicable in other forums.

A. "Banner"

Claims 18 and 34-41 recite a "banner." Patent Owner has acted as its own lexicographer and makes it clear that this term is to be construed "very broadly." Specifically, the '045 patent specification states:

> For purposes of the present invention, **the term 'banner' is meant to be construed very broadly** and includes any information displayed in conjunction with a web page wherein the information is not part of the same file as the web page. That is, a banner includes anything that is displayed or used in conjunction with a web page, but which can exist separately from the web page or which can be used in conjunction with many web pages. Banners can include graphics, textual information, video, audio, animation, and links to other computer sites, web sites, web pages, or banners. (GOOG 1001, 2:28-37.)

Under BRI, and given the explicit definition in the specification, a POSA would have understood the term "banner" to mean "information displayed in conjunction with a web page wherein the information is not part of the same file as the web page." This would include one or more of graphics, textual information, video, audio, animation, and links to other computer sites, web sites, web pages, or banners. (GOOG 1003, ¶ 39.)

B. "Best Suited"

Claims 7 and 9 recite a determining a server that is "best suited" to serve a banner. The '045 patent specification states:

> Typically, the information server best suited to handle the serving or transmittal of a banner to the terminal 36 will be the information server that can download or serve the banner to the terminal 36 in the shortest period of time. Other selection criteria can be used, however, in download or serve a banner to a terminal, including the network topological distance between the terminal 36 and the information servers, the geographical distance between the terminal 36 and the information servers, the bandwidth of the information servers, or the round trip times for a message between the terminal 36 and the information servers. (GOOG 1001, 20:62-21:7.)

Based on these examples, under BRI, a POSA would have understood the a "best suited"[4] server to at least include a server that can serve a banner based on one of the criteria including shortest period of time, network topological distance, geographical distance, bandwidth of the server, and round trip times. (GOOG 1003, ¶ 41.)

C. "Content General Request Signal"

Claim 4 recites a "content general request signal." The '045 patent specification states:

---

[4] The term "best suited" is relative and subjective, and is arguably indefinite under MPEP § 2173.05(b). Despite this defect, Petitioner makes a good-faith attempt to apply the claims.

> In other words, the initial banner request signal generated by terminal 36 during the step 112 can be a content general signal and may contain only the minimum amount of information needed to tell a designated computer site, information server, or other device which receives the initial banner request signal and on which a banner may or may not be stored or located, only that the terminal 36 desires that an unspecified banner be served to the terminal. (GOOG 1001, 15:8-16.)

Further, "[i]f the optional selection step 113 is used with the method 110, the terminal 36 will only request during step 112 that a banner be served to the terminal 36, but the terminal 36 will not specify which banner is to be served to the terminal 36." (GOOG 1001, 15:25-29.)

The '045 patent specification additionally states:

> A general content URL address for a banner does not provide the necessary information to determine which banner is to be displayed. Rather a general content URL address for a banner only indicates that a banner is to be displayed and the receiver of the signal generated by the terminal 36 during the step 112 can decide which banner is to be displayed during the selection step 113. A general content URL address for a banner could be of the form http://www.bannersite1.com/image;spacedesc=contentsitename.
> (GOOG 1001, 16:50-58.)

Further, "[t]he space descriptor field in the general content URL address can

reference different groups of banners such as, for example, a collection of car advertisements, a collection of detergent advertisements, etc., depending on the web page providing the general content URL address." (GOOG 1001, 17:3-8.) Accordingly, the content general request signal can still contain general information regarding a type of content or user interest, as long as a specific banner is not identified.

Under BRI, a POSA would therefore have understood the term "content general request signal" to mean "a request indicating that information is to be displayed and that the receiver can decide what information is to be displayed." (GOOG 1003, para 46.)

D. "Content Specific Request Signal"

Claim 5 recites a "content specific request signal." The '045 patent specification states: "If the optional selection step 113 is not used with the method 110, the terminal 36 will request during the step 112 that a specific banner to be served to the terminal 36." (GOOG 1001, 15:23-25.) Additionally:

> In order to speed up the process of downloading, transmitting, or serving a specific banner from an information server to the terminal 56, the content specific URL address of the requested or selected banner sent to the terminal during step 114 can contain the exact Internet Protocol (IP) address of the requested or selected banner. For example, instead of providing the specific content URL address for

> the banner 62 as http://www.bannersite1.com/banner1.gif, the specific content URL address for the banner 62 could be provided as, for example, http://236.45.78.190/banner1.gif, thereby removing any need to use the Domain Name System (DNS) to convert the alphanumeric address "www.bannersite1.com" of the information server to its exact IP address. (GOOG 1001, 18:62-19:8.)

Under BRI, a POSA would have understood the term "content specific request signal" to mean "a request containing a content specific URL address with the location of the information." (GOOG 1003, ¶ 49.)

VI. PERSON OF SKILL IN THE ART & STATE OF THE ART

A POSA is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity. With respect to the '045 patent, a POSA would typically have at least (a) a Bachelor of Science degree in computer science and/or a similar field or (b) at least 3 years of experience in web-based information management and delivery systems. (GOOG 1003, ¶¶ 13-14; GOOG 1005, ¶¶ 12-13.)

VII. IDENTIFICATION OF PRIOR ART AND CHALLENGE (37 C.F.R. § 42.104(b))

A. Prior art

(1) U.S. Patent No. 5,933,811 to Angles *et al.* was filed Aug. 20, 1996, and

issued Aug. 3, 1999. Angles is prior art under 35 U.S.C. 102(e)[5]. **(2)** U.S. Patent No. 5,948,061 to Merriman *et al.* was filed Oct. 29, 1996, and issued Sept. 7, 1999. Merriman is prior art under 35 U.S.C. 102(e). **(3)** U.S. Patent No. 5,960,409 to Wexler was filed Oct. 11, 1996, and issued Sept. 28, 1999. Wexler is prior art under 35 U.S.C. 102(e). **(4)** Fielding *et al.*, "HTTP Working Group Internet Draft Hypertext Transfer Protocol – HTTP/1.0" ("HTTP1.0") was published Feb. 20, 1996. HTTP1.0 is prior art under 35 U.S.C. 102(b). **(5)** Garland et al., "Implementing Distributed Server Groups for the World Wide Web" was published Jan. 25, 1995. Garland is prior art under 35 U.S.C. 102(b). **(6)** U.S. Patent No. 5,796,952 to Davis *et al.* was filed Mar. 21, 1997, and issued Aug. 18, 1998. Davis is prior art under 35 U.S.C. 102(e). **(7)** Meeker, Mary, "Technology: Internet/New Media The Internet Advertising Report" was published by Morgan Stanley, U.S. Investment Research in January 1996. Meeker is prior art under 102(a).

B. Challenge

IPR is requested for claims 1-12, 14-19, and 34-42 on the grounds for unpatentability listed in the index below. Per 37 C.F.R. § 42.6(d), copies of the references are filed herewith. In support of the proposed grounds, this Petition is accompanied by a declaration of a technical expert, Mr. Peter Kent (GOOG 1003),

---

[5] All references to Title 35 of the U.S. Code are to the pre-AIA version.

which explains what the art would have conveyed to a POSA.

Ground	35 USC	Index of Reference(s)	Claims
1	§103(a)	Angles, Merriman	**1**-6, 12, 14, 15, 17-19, **34**, 35, 40
2	§103(a)	Angles, Merriman, Garland	7-11, 16, 39
3	§103(a)	Angles, Merriman, Davis	42
4	§103(a)	Angles, Merriman, HTTP1.0	36-38, 41
5	§103(a)	Wexler, HTTP1.0	**1**-6, 12, 14-18, **34**-42
6	§103(a)	Wexler, HTTP1.0, Meeker	19
7	§103(a)	Wexler, HTTP1.0, Garland	7-11

1. Ground 1: Claims 1-6, 12, 14, 15, 17-19, 34, 35, and 40 Would Have Been Obvious Over Angles in view of Merriman

a) Independent Claim 1

Claim 1 does no more than add a few limitations to a conventional method of delivering "a second portion of information" (*e.g.*, a banner) referenced in "a first portion of information" (*e.g.*, a web page). The added limitations include a non-cache-blockable "first request signal" and a redirecting "location signal" providing an address of the second portion of information. These added limitations themselves were also well known in the art well before the '045 patent's filing date.

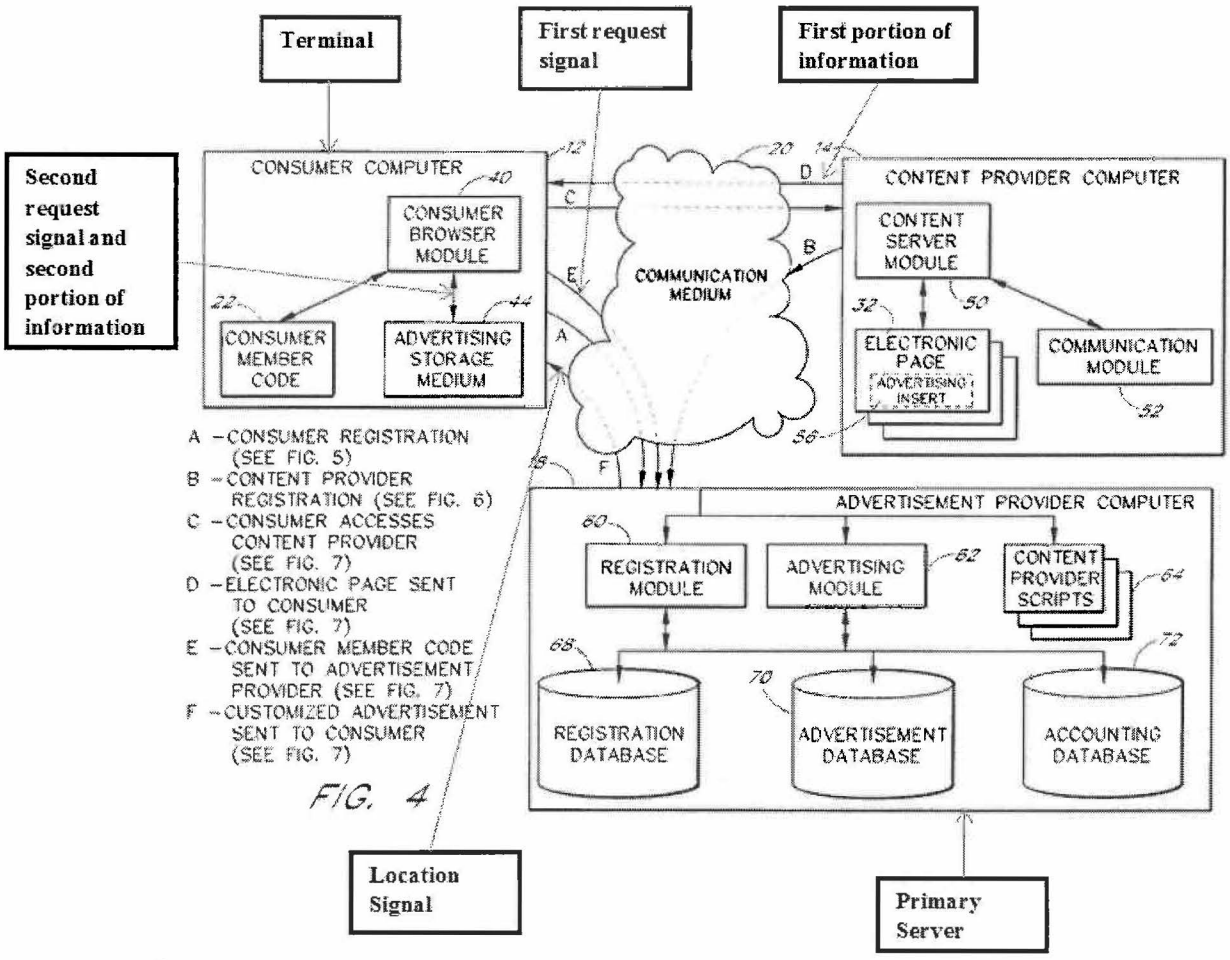

FIG. 4 of Angles (annotated above) illustrates the requesting and ultimate delivery of a customized advertisement to a consumer computer.

Angles' "Summary of the Invention" describes Angles' general process flow:

> "[T]he invention is directed to delivering custom advertisements to consumers who use their computers to view information offered by different content providers existing on the Internet. Preferably, when a consumer accesses a content provider, the content provider transmits an electronic document to the consumer. Embedded within the electronic document is a[n] advertisement request. When the

> consumer's computer displays the electronic document, the embedded advertisement request directs the consumer computer to communicate with an advertisement provider. In response, the advertisement provider provides a customized advertisement. The advertisement provider then tracks the consumer's response to the customized advertisement." (GOOG 1012, 2:59-3:5.) "The consumer computer the merges the content provider's electronic document with the advertisement provided by the advertisement provider to create a single displayed document to the consumer." (*Id.*, 3:58-65.)

The advertisement request in Angles is an unblockable signal. (GOOG 1003, ¶ 64.) In fact, Angles uses one of the exact same types of request that is described in the '045 patent – a CGI request. (*Id.*) This means that the signal is dynamic and the resulting content would not be cached. (*Id.*) Because the request identifies dynamic content that would not have previously been cached, a POSA would also have recognized that referencing a CGI script in the advertisement request prevents the advertisement request from being "cache-blocked" from reaching the advertisement provider computer 18. (*Id.*, ¶¶ 64-65.)

Instead of the advertisement provider actually providing a copy of the advertisement, the advertisement provider will provide a redirect signal that directs the consumer computer to another location. For example, "the consumer computer 12 receives an advertisement command which directs the consumer computer 12 to retrieve and display one of the advertisements stored on the advertising storage

medium 44." (*Id.*, 11:61-65.)

While the redirect signal from the advertisement computer is described in Angles as pointing to a location in local storage, it would have been obvious to a POSA that the redirect signal could just as easily point to a location on an external storage device, such as another server. (GOOG 1003, ¶¶ 68-69.) Angles delivers advertisement information using HTTP. (*See, e.g.*, GOOG 1012, 21:2-4.) A well-known feature of HTTP was the ability to send a redirect signal containing a URL address for another location where content could be found. Merriman discloses just such a signal. Specifically, Merriman teaches an HTTP redirect message containing an URL address of another server. (GOOG 1013, 7:24-26.)

Merriman discloses a system "for targeting the delivery of advertisements over a network." (GOOG 1013, Abstract.) Similar to Angles' customized advertisement, Merriman also provides solutions to "permit targeting of the advertisements of individual users." (*Id.*) Also similar to Angles, Merriman's system uses HTTP as the delivery mechanism. (*See, e.g.*, GOOG 1013, FIG. 1, block 14.) While Merriman discloses an overall banner advertisement system, it is referenced here primarily for its disclosure of an advertisement redirect. In Merriman, the redirect process is described as follows: "the user's browser again transmits a message to the ad server. The ad server notes the address of the computer of the browser…and transmits back the URL of the advertiser's web page

so that the user's web browser 16 generates a message 26 to contact the advertiser's web site." (GOOG 1013, 3:66-4:5; 7:22-26.) In response to receiving the redirect message, "the user's web browser 16 generates a message 26 to contact the advertiser's web site." (*Id.*, 4:4-5.) While Merriman describes its redirect methodology in response to a user's click, the fundamental utility of the redirect process is not limited thereto. It would have been obvious to a POSA that such redirect methodology to an external location would be a useful substitution for Angles' redirect to a local location. (GOOG 1003, ¶¶ 68-69.)

Accordingly, it would have been obvious to a POSA to replace Angles' advertising command identifying a location of an advertisement on local storage with Merriman's HTTP redirect message identifying a location of an advertisement on networked storage. Such replacement would have been nothing more than a simple substitution of one known element (Merriman's HTTP redirect message) for another (Angles' advertisement command) to obtain predictable results (an advertisement command from an advertisement server that identifies a URL at which the advertisement can be obtained). (GOOG 1003, ¶ 68; *see*, *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 416 (2007).) Also, a POSA would have been motivated to replace Angles' advertisement command with Merriman's HTTP redirect message for scalability reasons. (GOOG 1003, ¶ 31.)

Further specifics regarding the applicability of Angles in view of Merriman

to claim 1 are provided in the below claim chart:

Claim 1 limitation	Angles/Merriman disclosures
[1.P] A method for storing information on a primary server and one or more secondary servers and on computer sites connected to a computer network, wherein information delivered over the computer network to a terminal or a group of terminals may contain references to other information to be delivered to the terminal, comprising	In Angles, the electronic document is the "information delivered over the network to a terminal," the advertisement is the "other information," and the advertisement request embedded in the electronic document is the "reference[] to other information." (GOOG 1003, ¶ 62.)
[1.1] serving a first portion of information to a terminal, wherein said first portion of information contains a reference to a second portion of information;	In Angles, "[p]referably, when a consumer accesses a content provider, the content provider transmits an electronic document to the consumer." (GOOG 1012, 2:62-64.) This constitutes "serving" the electronic document. (GOOG 1003, ¶ 62.) "Embedded within the electronic document is a[n] advertisement request." (GOOG 1012, 2:64-3:5.) The "electronic document" is the recited first portion of information. (GOOG 1003, ¶ 24.) The "customized advertisement" from the advertisement provider is the recited second portion of information. (*Id.*) The "embedded advertisement request" is the recited reference to the second portion of information. (*Id.*)
[1.2a] causing a first request signal to be transmitted from the terminal to a primary server requesting a location address for said second portion of	In Angles, "the embedded advertisement request directs the consumer computer to communicate with an advertisement provider." (GOOG 1012, 2:67-3:2.) As a result, "an advertising request is sent to the advertisement provider computer." (*Id.*, Abstract.) "The advertising request requests a location of the customized advertisement, so that the advertisement provider

information from which said second portion of information can be served to the terminal,	computer returns an advertisement command indicating a location of the advertisement." (GOOG 1003, ¶ 63.) Further, "the advertisement command identifies a particular location on the advertising storage medium 44 … where an advertisement is located." (GOOG 1012, 11:66-12:2.) The "particular location" in Angles is the "location address" of claim 1.
[1.2b] wherein said first request signal includes information intended to prevent said first request signal from being blocked from reaching said primary server by either the terminal or any intermediary device located topologically between the terminal and the primary server as a result of previous caching of said first portion of information or said second portion of information in the terminal or said intermediary device;	In Angles, "advertisement request 26 references a content provider CGI script 64 which exists on the advertisement provider computer 18." (GOOG 1012, 13:2-4; see also 7:65-8:1.) A POSA would have recognized that a server (such as the advertisement provider computer in Angles) executes a requested CGI script to dynamically generate content (*e.g.*, a customized advertisement). (GOOG 1003, ¶ 64.)
[1.3] sending a location signal from the primary server to the terminal providing said location address of said second portion of information;	In Angles, "the advertising module 62 in the advertisement provider computer obtains the appropriate advertisement command from the advertisement database…. The advertising module then sends the advertisement command to the consumer computer 12." (GOOG 1012, 23:37-39.) The "advertisement command identifies **a particular location** on the advertising storage medium 44, such as the particular track and sector where an advertisement is located." (*Id.*, 11:66-12:2.) Because the advertisement command identifies a particular location, it constitutes the "location address" as recited in claim 1.

	In Merriman, the advertisement server "sends the redirect message causing the user's browser to receive the URL for the advertiser's web site based upon data stored in the server." (*Id.*, 7:22-26.) In response to receiving the redirect message, "the user's web browser 16 generates a message 26 to contact the advertiser's web site." (*Id.*, 4:4-5.)
[1.4] causing a second request signal to be transmitted from the terminal containing said location address of said second portion of information and requesting said second portion of information be served to the terminal; and	In Angles, "the consumer computer 12 receives an advertisement command which directs the consumer computer 12 to retrieve and display one of the advertisements stored on the advertising storage medium 44." (GOOG 1012, 23:32-35; GOOG 1003, ¶ 71.) In the Angles-Merriman combination, the advertisement command received by the consumer computer is an HTTP redirect signal identifying the URL at which the advertisement can be found and requesting return of the advertisement. (GOOG 1003, ¶ 71.) In Merriman, "the user's web browser 16 generates a message 26 to contact the advertiser's web site" in response to the HTTP redirect message. (GOOG 1013, 4:2-5.) It would have been obvious to a POSA that Angles in view of Merriman further discloses sending an HTTP redirect message to the consumer computer to retrieve the customized advertisement from another computer. (GOOG 1003, ¶ 71.)
[1.5] serving said second portion of information to the terminal.	In Angles, "the consumer computer 12 receives an advertisement command which directs the consumer computer 12 to retrieve and display one of the advertisements stored on the advertising storage medium 44." (GOOG 1012, 23:32-35.) In the Angles-Merriman combination, the advertisement command directs the consumer computer to retrieve and display one of the advertisements stored at the location identified by the URL in the HTTP redirect message. (GOOG 1003, ¶ 73.) As a result, the advertisement is served to the consumer computer for display. (*Id.*)

For the reasons provided above, Angles in view of Merriman renders claim 1 obvious.

b) Independent claim 34

Claim 34 recites similar features to claim 1, and is unpatentable for the reasons discussed with respect to claim 1. The differences between claim 34 and claim 1 are also addressed here.

Instead of a "first portion of information," claim 34 recites "a web page." It would have been obvious to a POSA that Angles' electronic document constitutes a web page. (GOOG 1003, ¶ 75.) Instead of a "terminal," claim 34 recites a "computer." Angles' consumer computer meets this limitation. (GOOG 1012, 10:20-42.) Instead of a "second portion of information," claim 34 recites a "specific banner." This too is disclosed in Angles, which discloses a customized advertisement 30 to be merged into electronic page 32. (GOOG 1003, ¶ 76.)

Claim 34 also recites, "determining which specified banner will be served to the computer." The customized advertisement of Angles, presented in a web page and contains a hyperlink to more advertising information, is a specified banner. (GOOG 1003, ¶¶ 76-77.) It is clear in Angles that "the advertisement computer 18 **selects** an appropriate customized advertisement 30." (GOOG 1012, 8:13-15)

Instead of a "first request signal," claim 34 recites a "banner request signal" that "includes a Uniform Resource Locator address for said primary server."

Angles' advertising request signal constitutes a banner request signal, as the requested customized advertisement is a banner. Angles' advertisement request signal also contains a URL. (GOOG 1012, 13:4-7.)

Lastly, instead of a "location signal…providing said location address" of the advertisement, claim 34 recites, "said banner location signal includes the Uniform Resource Locator address for a device on which the specific banner to be served to the computer is stored." As discussed with respect to claim 1, Angles' advertising command constitutes the recited banner location signal. (GOOG 1012, 23:37-39, 11:66-12:2.) While Angles does not disclose that the advertising command contains a URL, as discussed with respect to claim 1, replacing the reference to local storage in Angles' advertising command with a reference to networked storage as described in Merriman would have been obvious to a POSA. (GOOG 1003, ¶¶ 67-68.) As described in Merriman, the HTTP redirect message contains a URL. (GOOG 1013, 7:22-26.) As a result, Angles in view of Merriman renders obvious an HTTP redirect message that redirects the consumer computer to send a request signal to retrieve a customized advertisement (*i.e.*, "banner") from another server. (GOOG 1003, ¶¶ 67-68.)

For the reasons provided above and for those provided with respect to claim 1, Angles in view of Merriman would have rendered claim 34 obvious.

c) Dependent Claims 2-6, 12, 14, 15, 17-19, 35, and 40

Claim 2 depends from claim 1 and recites "wherein said intermediary device is a proxy server." A POSA would have understood that an intermediary device may exist topologically between any two networked components in Angles, such as the advertising computer and the consumer computer. (GOOG 1003, ¶ 80.) Indeed, Angles discloses that the consumer computer may need to connect through an Internet Service Provider before connecting to the Internet (GOOG 1012, 9:35-44; 9:56-64; Fig. 2, element 34) or that "the consumer computers may be connected to a local area network which in turn is directly connected to the Internet." (GOOG 1012, 9:47-49.) Angles further discloses the use of HTTP as the "standard World Wide Web client-server protocol used for the exchange of information (such as HTML documents, and client requests for such documents) between a Web browser and a Web server." (*Id.*, 6:49-53.)

While Angles does not explicitly disclose proxy servers as the intermediary devices, proxy servers were well known intermediaries in systems using HTTP protocols. (GOOG 1003, ¶ 80.) For example, Merriman discloses using a proxy server: "[i]ncluded in each message 23 typically to the advertising server 19 are: ... the operating system of the computer on which the browser is operating and the proxy server type." (GOOG 1013, 3:44-52.) It would have been obvious for a POSA to use the proxy server disclosed in Merriman as Angles' intermediary,

because the combination would have made use of standard network topology (GOOG 1003, ¶ 80.)

Claim 3 depends from claim 2 and recites "wherein said second portion of information is served from the primary or secondary servers." Also, claim 6 depends from claim 1 and recites, "wherein said second portion of information is served from one of the secondary servers." As discussed above with respect to claim 1, the Angles-Merriman combination discloses that a server other than the advertisement provider sends the customized advertisement to the consumer computer in response to receiving an HTTP redirect message. (GOOG 1003, ¶¶ 84 and 95.) This other server constitutes a secondary server. (*Id.*)

Claim 4 depends from claim 1 and recites "wherein said first request signal is a content general request signal." Claim 35 depends from claim 34 and recites, "wherein said banner request signal includes a content general Uniform Resource Locator address." As discussed above, under BRI, a content general request signal is simply a request indicating that information is to be displayed and that the receiver can decide what information is to be displayed. As discussed above with respect to claim 1, Angles teaches the first request signal in the form of the advertising request. The advertising request does not specify which customized advertisement is to be served; instead, the advertisement provider computer determines which customized advertisement to server after receiving the

advertising request. (GOOG 1012, Abstract; GOOG 1003, ¶ 87.) The advertising request contains "the Internet address or URL of the advertisement provider computer 18." (GOOG 1012, 13:4-7.) Accordingly, the advertising request of Angles discloses the content general request signal as defined by the '045 patent. (GOOG 1003, ¶¶ 87-88.)

Claim 5 depends from claim 4 and recites "wherein said second request signal is a content specific request signal." Claim 40 depends from claim 35 and recites, "wherein said banner location signal includes a content specific Uniform Resource Locator address for the specified banner." As discussed above with respect to claim 1, Angles discloses that "a short advertisement command can be sent [to the consumer computer] which **specifically** retrieves a **particular advertisement** from the advertising storage medium." (GOOG 1012, 23:52-54.) To the extent that Patent Owner may argue retrieval from the advertising storage medium does not constitute a signal, as discussed with respect to claim 1, Angles in view of Merriman renders obvious sending an HTTP request including a URL to retrieve the specific advertisement from another server. (GOOG 1003, ¶¶ 91-92.)

Claim 12 depends from claim 1 and recites "storing said second portion of information in the terminal." In Angles, the "consumer computer 12 stores the customized advertisements 30 on the advertising storage medium 44." (Angles, 23:28-29.) The advertising storage medium is local to the consumer computer.

(Angles, 12:50-65; GOOG 1003, ¶ 98.)

Claim 14 depends from claim 4 and recites "including selecting the composition of said second portion of information."[6] In Angles, "[b]ased on the consumer's profile, the advertisement provider computer 18 **selects** an appropriate customized advertisement 30. (GOOG 1012, 8:13-17.) This constitutes "selecting the composition" of the advertisement. (GOOG 1003, ¶ 101.)

Claim 15 depends from claim 14 and recites "wherein the results of said composition selection are included in said location signal sent from the information server to the terminal."[7] Angles' advertisement command "directs the consumer control module 42 to retrieve the customized advertisement 30 from the advertising storage medium" after the customized advertisement has been selected by the advertising provider computer. (GOOG 1012, 23:32-35.) Since identification of the customized advertisement is the "result[] of said composition selection," claim 15 would have been obvious over Angles in view of Merriman (GOOG 1003, ¶ 104.)

Claim 17 depends from claim 1 and recites "wherein said reference to a second portion of information includes at least a portion of a URL." As discussed

---

[6] Petitioner notes that the claimed "composition" has no antecedent basis, and the '045 patent disclosure provides no discussion related to any composition. Despite this defect, Petitioner makes a good-faith attempt to analyze the claims.

[7] See fn. 6.

above, Angles' advertisement request is the reference to the customized advertisement, and contains a URL. (GOOG 1003, ¶ 107.) Specifically, Angles discloses: "the advertising request 26 is an HTML tag which identifies 1) the content provider script and 2) the content provider member code and 3) the Internet address or **URL** of the advertisement provider computer 18." (GOOG 1012, 13:4-7.)

Claim 18 depends from claim 1 and recites "wherein said first portion of information is a web page, said second portion of information is a banner, and said reference is a link." As discussed, Angles' electronic document is the first portion of information. (GOOG 1003, ¶ 75.) A POSA would have recognized that an HTML document served by a web server and transferred using the HTTP protocol is a web page. (*Id.*) As also discussed, Angles' customized advertisement is the second portion of information. (*Id.*) The customized advertisement is an "advertisement insert" contained in the electronic web page. "The advertisement insert 56 is a place-holder configured to contain the customized advertisement 30." (GOOG 1012, 12:51-58.) Further, "the customized advertisement 30 contains a hyper-link to more advertising information." (*Id.*, 20:5-7.) Customized advertisement 30 thus constitutes a banner. (GOOG 1001, 2:35-37.)

And, such reference is a link. (GOOG 1003, ¶¶ 76 and 111.) As Angles explains, "advertisement request 26 is an HTML tag which directs the consumer

computer 12 to establish a communication link with the advertisement provider computer 18....[and] the HTML tag in the advertisement request 26 directs the advertising module 62 to execute the content provider script 64." (GOOG 1012, 15:6-11.)

Claim 19 depends from claim 1 and recites "including counting at least one display of said second portion of information on the terminal." Angles' advertisement provider computer counts a number of displays after receiving the advertisement request and extracting the content provider member code. (GOOG 1003, ¶ 114.) "The advertisement provider uses the content provider member code to **track the number of advertisements displayed** by a particular content provider." (GOOG 1012, 3:46-53.) This counting method described in Angles—counting the number of non-blockable requests as the number of displays—is also how the '045 patent counts each display of its banners. (GOOG 1001, 15:30-35.)

Accordingly, claims 2-6, 12, 14, 15, 17-19, 35, and 40 would have been obvious over Angles in view of Merriman.

2. Ground 2: Claims 7-11, 16, and 39 would have been obvious over Angles in view of Merriman and further in view of Garland

Angles in view of Merriman renders obvious delivering advertisement information through the use of Web pages using multiple servers and the HTTP protocol. As discussed in Merriman, multiple advertisement servers may be used.

(GOOG 1013, 4:20-21; GOOG 1003, ¶ 120.) While neither Angles nor Merriman discusses how to choose between multiple servers, such server selection methodologies were well known in the art. For example, Garland discloses the concept of load balancing "to implement a Distributed Web server which will spread incoming request load among several machines" with the use of HTTP. (GOOG 1009, p. 1.)

It would have been obvious to a POSA to incorporate the load balancing of web resources as taught by Garland with the delivery of web-based advertising banners using HTTP as disclosed by Angles in view of Merriman, so as to manage the multiple ad servers disclosed in the Angles-Merriman combination. (GOOG 1003, ¶ 120.) Such a combination would have been obvious for a skilled artisan as it is nothing more than applying a known technique (control of web resources and requests as taught by Garland) to a known method (the multi-server, web-based banner advertising delivery process of Angles-Merriman) to yield predictable results (the advertisement provider computer selecting a least loaded (or best suited) advertisement provider computer to serve a banner). (*Id.*) A POSA would also have been motivated to modify the Angles-Merriman combination by selecting the least loaded web server in Garland to further address the transmission delay issue as addressed in Angles. (*Id.*)

Claim 7 depends from claim 1 and recites "wherein after the primary server

receives the first request signal from the terminal, further including determining which server connected to the computer network is best suited for serving said second portion of information to the terminal." Claim 7 is rendered obvious by the combination of Angles, Merriman, and Garland. (GOOG 1003, ¶ 118.) While the combination of Angles and Merriman is discussed above with respect to claim 1, Angles and Merriman do not explicitly disclose selecting a "best suited" server. But it was well known in the art to select the least loaded server from a group of servers. (*Id.*) For example, as discussed above, Garland teaches that "[i]ncoming requests are received by the dispatcher, and they are redirected to one of the member servers for actual processing. It is also the task of the dispatcher to balance the load placed upon the various members under its control. In essence, it tries to always dispatch incoming requests to **the least loaded member**." (GOOG 1009, p. 3.) A POSA would have recognized that selecting the least loaded server would include selecting which a server can serve the client in the shortest period of time. (GOOG 1003, ¶ 118.) Selecting a server based on service in the shortest period of time constitutes selecting the "best suited" server. (See Section V above.)

Both Angles and Merriman deliver advertisement information using HTTP. (GOOG 1003, ¶ 119.) Garland discloses the concept of load balancing in an HTTP environment "to implement a Distributed Web server which will spread incoming request load among several machines." (GOOG 1009, p. 1.) As discussed above, it

would have been obvious to POSA to incorporate the load balancing of Web resources as taught by Garland with the delivery of web-based electronic advertising as taught by Angles in view of Merriman. (GOOG 1003, ¶ 120.)

Claim 8 depends from claim 7 and recites "wherein results of said determining are included in said location signal sent from the information server to the terminal." As discussed above, Garland discloses selecting the least loaded server. Garland further teaches that the HTTP redirect message "instructs a client that the requested object has been found on another server, **whose address is included in the body of the response**." (GOOG 1003, ¶ 123; GOOG 1009, p. 4.)

Claim 9 depends on claim 8 and recites "creating a matrix of selections between each of the terminals or groups of terminals and each of the servers and using said matrix to determine which of the servers is best suited to serve said second portion of information to the terminals or groups of terminals." Garland creates a matrix to test "each possible combination of server count, client count, and file access pattern" for latency optimization. (GOOG 1009, p. 6; GOOG 1003, ¶ 126.) Although Garland does not explicitly state that the measured results for each possible combination are used for selecting the "best suited" server, a POSA would have been motivated to use the measured results (including the measured latency) to select a server with the least latency according to a scheduling policy. (GOOG 1003, ¶ 127.) Garland wants "to distribute requests in such a way as to

balance the load placed on the various member servers. This requires two things: the dispatcher must have some knowledge of the load on members, and there must be some scheduling policy used to select which member to dispatch the next request to." (GOOG 1009, p. 5.) A POSA would have recognized that the measured latency for each possible combination would be used as a scheduling policy for load balancing as such scheduling policy further helps to "reduce[s] the latency of servicing a request." (GOOG 1009, p. 1; GOOG 1003, ¶ 127.)

Claim 10 depends from claim 9 and recites "wherein said selections contain round trip times between the servers and the terminals or groups of terminals." As discussed above, Garland teaches measuring latency for each possible combination of servers and clients. The measured latency includes the round trip time for each possible combination of servers and clients. (GOOG 1003, ¶ 130.)

Claim 11 depends from claim 1 and recites "including making one of the secondary servers a new primary server if the original primary server becomes inaccessible." It was well known in the art to use a group of web servers for handling failover situation such that functions of one web server are switched to a second server. (GOOG 1003, ¶ 133.) For example, Garland describes "solv[ing] their problem of server failure" through a round-robin resolution. (GOOG 1009, p. 8.) A POSA would have understood that using a group of web servers to solve server failure problem means that if the first primary server fails, a second server

can be made the new primary server. (GOOG 1003, ¶ 133.)

Claim 16 depends from claim 1 and recites "wherein said location signal includes an HTTP 302 redirect command." Claim 39 depends from claim 34 and recites, "wherein said banner location signal constitutes an HTTP 302 redirect signal." As discussed above, Angles in view of Merriman renders obvious sending an HTTP redirect message from one provider computer to the consumer computer to retrieve the customized advertisement from another provider computer. A POSA would have understood that such an HTTP redirect message would be an HTTP 302 redirect message. (GOOG 1003, ¶¶ 123 and 136.) For example, Garland states that "[i]n order to make the server group work, there needs to be some mechanism for the dispatcher to redirect a client's request to another server.... The status code of particular interest to us is: **Found (302)**." (GOOG 1009, p. 4.)

Accordingly, claims 7-11, 16, and 39 would have been obvious over Angles in view of Merriman and further in view of Garland.

3. Ground 3: Claim 42 would have been obvious over Angles in view of Merriman and further in view of Davis

Claim 42 depends from claim 34 and recites "wherein said device is said primary server." As discussed above, Angles in view of Merriman renders obvious an advertisement provider computer sending an HTTP redirect message to the consumer computer to retrieve the customized advertisement from another

advertisement provider computer. It would have been obvious to a POSA that the computer from which the customized advertisement is ultimately retrieved is simply a design choice. (GOOG 1003, ¶ 141.) A POSA would have recognized that, if the first advertisement provider computer stores the same advertisements as the second ad provider computer, the first advertisement provider computer would simply direct the consumer computer to retrieve the customized advertisement from that location on the same first advertisement provider computer. (*Id.*)

Such a limitation is described in Davis:

> "[w]hen a client machine passes a TCP/IP request for the Web page to the first server, the Web page is downloaded to the client, including the ad banner embedded using the <IMG> tag. The <IMG> tag is used to reference a resource (*i.e.*, the "ad banner") stored **on the same or a different server** which captures the user's ID (via the HTTP request header) and dynamically returns an ad related image to the client for display within the Web page." (GOOG 1014, 3:35-42.)

Combining Angles, Merriman, and Davis is simply combining prior art elements (storing banners on the same or a different server in Davis, and the ad provider computer that redirects the client to the banner location in Angles-Merriman) according to known methods to yield predictable results (an ad provider computer that directs the client to retrieve banners from the same ad provider computer). (GOOG 1003, ¶¶ 142.) As such, claim 42 would have been obvious

over Angles in view of Merriman and further in view of Davis.

4. Ground 4: Claims 36-38 and 41 would have been obvious over Angles in view of Merriman and further in view of HTTP1.0

Both Angles and Merriman disclose the use of HTTP for advertisement delivery. (GOOG 1003, ¶ 148.) At the time the '045 patent was filed, HTTP was the standard protocol for information delivery over the web, so a POSA would have been familiar with the standard functions of HTTP. (GOOG 1003, ¶¶ 18 and 20-21.) As such, it would have been obvious to a POSA to implement the functions of HTTP1.0 into the Angles-Merriman combination. (GOOG 1003, ¶ 147.)

Claim 36 depends from claim 34 and recites "determining whether said specified banner is stored on the computer." Claim 37 depends from claim 36 and recites "wherein after said determining whether said specified banner is stored on the computer, if said specified banner is not stored on the computer then including causing a second banner request signal to be sent to said device requesting that said device serve said specified banner to the computer." But such limitations simply recite the default behavior of sending an HTTP request. (GOOG 1003, ¶¶ 146-147.) HTTP1.0 states that any client or server may serve as a cache. (GOOG 1008, p. 6.) If a participant does not have a cached response, then any signal must be sent on to its destination. (GOOG 1003, ¶ 148.)

Claim 38 depends from claim 37 and recites "serving the specified banner

from said device to said computer." As discussed with respect to claim 1, Angles serves a customized advertisement to a consumer computer. (GOOG 1003, ¶ 151.)

Claim 41 depends from claim 34 and recites "tagging said specified banner as being cachable." Again, this is just standard functionality of HTTP. As described in HTTP1.0, "Expires entity-header field gives the date/time after which the entity should be considered stale. This allows information providers to suggest the volatility of the resource, or a date after which the information may no longer be val*Id.* Applications must not cache this entity beyond the date given." (GOOG 1008, p. 33.) That is, it would have been obvious to a POSA that a web browser application would cache a response (such as a banner) based on this tag if it is before the date given in the Expires entity-header field. (GOOG 1003, ¶ 154.) It would have been obvious to incorporate such standard HTTP functionality into the Angles-Merriman combination.

Claims 36-38 and 41 thus would have been obvious over Angles in view of Merriman and further in view of HTTP1.0.

5. Ground 5: Claims 1-6, 12, 14-18, and 34-42 would have been obvious over Wexler in view of HTTP1.0

a) Independent Claim 1

Wexler discloses a "system and method for providing on-line third party accounting and statistical information" where "a banner, displayed for the purpose

of enticing a first party (user) to visit a fourth party's (advertiser) Web site, is served to the user's Web browser by a second party (banner published)." (Wexler, Abstract.) Wexler further discloses the use of "[a] Hypertext Transfer Protocol (HTTP) server program." (Wexler, 5:15-16.)

Figure 2 of Wexler (annotated below to correlate to claim 1) illustrates the requesting and ultimate delivery of an advertiser's page.

FIG. 2

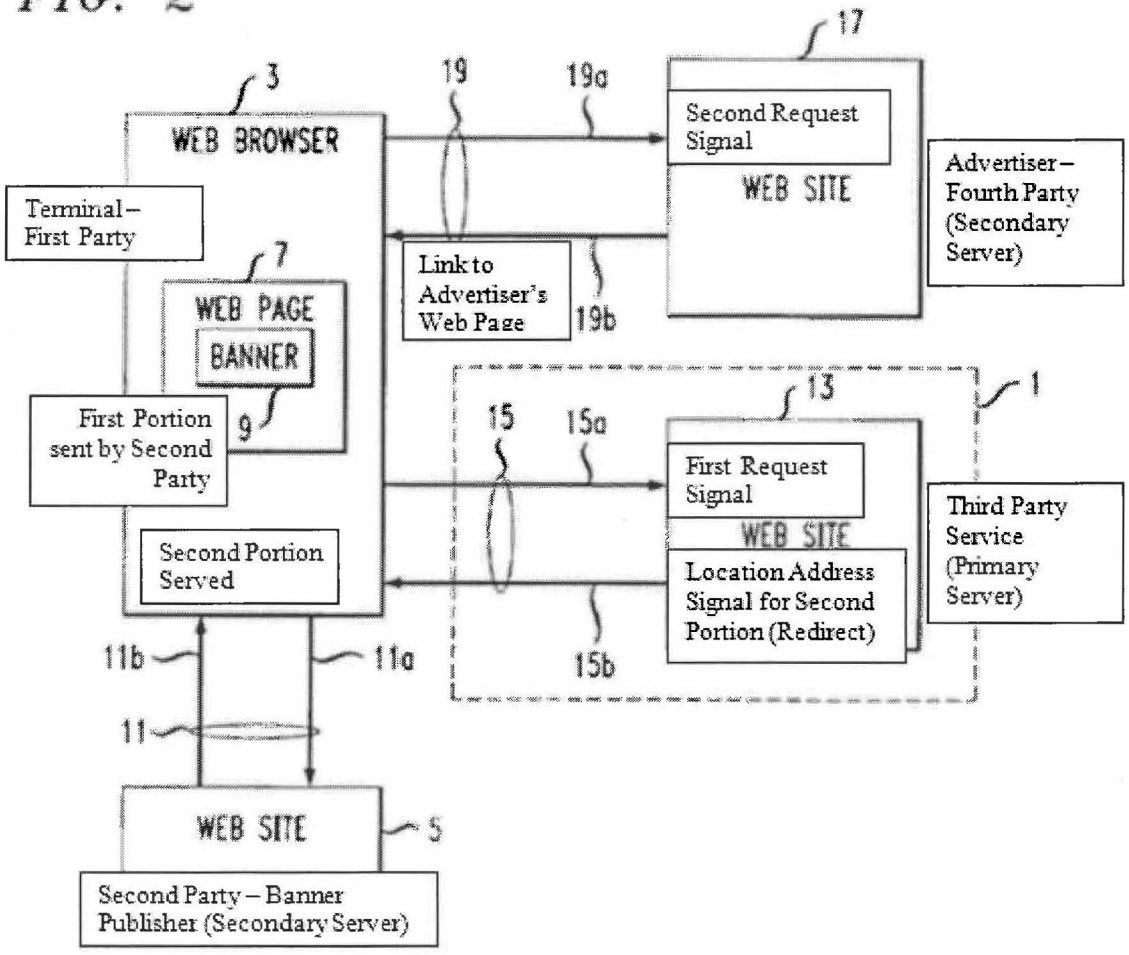

In claim 1 of the '045 patent, the returned advertisement is referred to as the "second portion of information." In claim 34 of the '045 patent, the returned

advertisement is referred to as the "banner." The advertiser's web page in Wexler, which is returned by Wexler's redirect process, reads on the claimed "banner" under the BRI, discussed above. (GOOG 1003, ¶ 158.) The advertiser's web page is "information displayed in conjunction with another web page." (*Id.*) The advertiser's web page is not part of the same file as web page 7, as the advertiser's web page is from web site 17 and web page 7 is from web site 5. (*Id.*) It would have been obvious to a POSA that Wexler's advertiser's web page could be shown in the main screen of the browser, in a different browser tab or window, in a frame or an iframe within another web page, or in a pop-up window. (*Id.*) Accordingly, the advertiser's web page in Wexler constitutes the claimed "banner."

For the purposes of comparison to the claims of the '045 patent, Wexler's process starts with web page 7 and banner 9 already assembled, such that banner 9 forms "a hypertext link to the third party Web site 13" where "the third party accounting and statistical service 13 receives a download request signal 15a from the user's Web browser 3" and the "third party service 13 accepts the download request signal 15a and increments a counter that keeps track of the number of received request signals." (GOOG 1007, 3:50-60.) The "third party service 13 redirects the received download request signal to the advertiser's Web site 17" where "the redirect request 15b is sent to the user's Web browser 3 from the third party Web site 13, and, from the browser, a download request 19a is sent to the

advertiser's Web site 17." (*Id.*, 5:3-8.) "Once the download request signal is received by the advertiser, the advertiser's Web page is downloaded to the user's Web browser 3." (*Id.*, 5:9-11.)

While Wexler discloses a counter to keep track of the number of times a banner is displayed at the client device, Wexler does not disclose that the request signal 15a is "unblockable." But it was well known in the art at the time to employ a "cache-busting" mechanism to avoid caches for the purpose of accurately counting banner advertisements. (GOOG 1003, ¶¶ 15 and 17.) For example, HTTP1.0 discloses using a Pragma no-cache directive to prevent caching. (*Id.*, ¶ 162.) This is one of the same types of signals referred to in the '045 patent specification as "unblockable." (GOOG 1001, 18:14-17.)

Wexler refers to hypertext throughout its specification, and specifically suggests the use of HTTP as its transmission protocol. (GOOG 1003, ¶ 163.) HTTP was the standard protocol for information delivery over the web at the time, so a POSA would have been familiar with the standard functions of HTTP, as described in HTTP1.0. (*Id.*) Accordingly, it would have been obvious to a POSA to incorporate standard HTTP functionality as set forth in HTTP1.0 into the HTTP-based advertising system of Wexler. (*Id.*) Further, using HTTP1.0 constructs to implement the advertising system of Wexler would simply have been a combination of prior art elements (an HTTP request in Wexler and the Pragma

header in HTTP1.0) according to known methods to yield predictable results (an HTTP request modified with the Pragma header). (*Id.*) Accordingly, it would have been obvious to a POSA to implement the first request signal (signal 15a) in Wexler's hypertext-based advertising system using the Pragma no-cache directive as set forth in HTTP1.0. (GOOG 1003, ¶ 163.)

Further specifics regarding the applicability of Wexler in view of HTTP1.0 to claim 1 are provided in the below claim chart:

Claim 1 limitation	Wexler/HTTP1.0 disclosures
[1.P] A method for storing information on a primary server and one or more secondary servers and on computer sites connected to a computer network, wherein information delivered over the computer network to a terminal or a group of terminals may contain references to other information to be delivered to the terminal, comprising	For purposes of analysis here, the "information delivered over the computer network" is the web page 7 assembled together with the banner 9 of Wexler. (GOOG 1003, ¶¶ 157 and 159.) The "reference[] to other information to be delivered to the terminal" is the "hypertext link to the third party Web site 13" included in the displayed banner 9. (*Id.*)
[1.1] serving a first portion of information to a terminal, wherein said first portion of information contains a reference to a second portion of information;	The service of banner 9 together with web page 7 to the user's browser constitutes "serving a first portion of information to a terminal." (*Id.*, ¶ 159.) The "reference to a second portion of information" is Wexler's hyperlink to the third party's web page requesting an advertiser's web page, while the advertiser's web page is the second portion of information. (*Id.*)
[1.2a] causing a first	The "first request signal" in Wexler is signal 15a,

request signal to be transmitted from the terminal to a primary server requesting a location address for said second portion of information from which said second portion of information can be served to the terminal,	transmitted from the user's terminal by a web browser 3. "If a user clicks on the banner 9 forming a link to the third party, then, as indicated in operation block 103 of FIG. 3, the third party accounting and statistical service 13 [primary server] receives a download request signal 15a [first request signal] from the user's Web browser 3." (GOOG 1007, 4:54-57; GOOG 1003, ¶ 160.)
[1.2b] wherein said first request signal includes information intended to prevent said first request signal from being blocked from reaching said primary server by either the terminal or any intermediary device located topologically between the terminal and the primary server as a result of previous caching of said first portion of information or said second portion of information in the terminal or said intermediary device;	HTTP1.0 discloses the use of a Pragma No-Cache directive to avoid caching a request: "The Pragma general-header field is used to include implementation specific directives that may apply to any recipient along the request/response chain. All pragma directives specify optional behavior from the viewpoint of the protocol; however, some systems may require that behavior be consistent with the directives. When the _no-cache_ directive is present in a request message, an application should forward the request toward the origin server even if it has a cached copy of what is being requested. This allows a client to insist upon receiving an authoritative response to its request. It also allows a client to refresh a cached copy which is known to be corrupted or stale. "Pragma directives must be passed through by a proxy or gateway application, regardless of their significance to that application, since the directives may be applicable to all recipients along the request/response chain." (GOOG 1008, pp. 35-36.)
[1.3] sending a location signal from the primary server to the terminal providing said location address of said second portion of information;	In Wexler, redirect request 15b is the recited location signal. "The download request received by the third party service 13 is ultimately intended to obtain information from the advertiser. As such, the third party service 13 redirects the received download request signal to the advertiser's Web site 17, as indicated in operation block 107. Specifically, the redirect request

	15b [location address signal] is sent to the user's Web browser 3 from the third party Web site 13." (GOOG 1007, 5:1-8.)
[1.4] causing a second request signal to be transmitted from the terminal containing said location address of said second portion of information and requesting said second portion of information be served to the terminal; and	The "second request signal" in Wexler is signal 19a sent by the terminal to the advertiser's server (secondary server) to have the secondary server send the second portion of information to the terminal, *i.e.*, the advertiser's web page 19b. (GOOG 1003, ¶ 165.)
[1.5] serving said second portion of information to the terminal.	In Wexler, the browser's downloading the linked advertiser's web site constitutes "serving the second portion of information." (GOOG 1003, ¶ 165.)

For the above reasons, claim 1 would have been obvious over Wexler in view of HTTP1.0.

b) Independent claim 34

Claim 34 recites similar features to claim 1, and is unpatentable for the reasons discussed with respect to claim 1. The differences between claim 34 and claim 1 are addressed here.

Instead of a "first portion of information," claim 34 recites "a web page." Wexler's web page embedded with a banner ad is the "first portion of information." (GOOG 1003, ¶ 167.) Instead of a "terminal," claim 34 recites a "computer." Wexler's user's computer running web browser 3 meets this limitation. (GOOG

1007, 3:37-39.) Instead of a "second portion of information," claim 34 recites a "specific banner." As discussed above, the advertiser's web page 17 returned as a result of redirect signal 19a is the recited "specific banner." (GOOG 1003, ¶¶ 158 and 167.)

Claim 34 also recites, "determining which specified banner will be served to the computer." A POSA would have understood that the third party service 13 in Wexler first needs to decide which advertiser's web page to redirect the user's web browser to before the third party service 13 sends the redirect request 15b to the user's web browser. (GOOG 1003, ¶ 168.)

Instead of a "first request signal," claim 34 recites a "banner request signal" that "includes a Uniform Resource Locator address for said primary server." As discussed above, signal 15a of Wexler is the banner request signal, which includes the URL of third party service 13. (GOOG 1007, 4:41-57.) As described in HTTP1.0, request signals typically include the Uniform Resource Identifier of the location from which information is requested (here the primary server), also known as a Uniform Resource Locator. (GOOG 1008, pp. 11 and 19-20.)

Lastly, instead of a "location signal…providing said location address" of the advertisement, claim 34 recites, "said banner location signal includes the Uniform Resource Locator address for a device on which the specific banner to be served to the computer is stored." In Wexler the "device" is the advertiser's Web site 17.

"When the specific URL is requested, the request itself and Web browser 3 information is recorded, and **the redirect to the intended URL**, *i.e.*, the advertiser's Web site, is issued. In this manner, the user is transparently redirected to the advertiser's Web site 17." (GOOG 1007, 5:18-23.) As described in HTTP1.0, an HTTP redirect message includes the new URL. (GOOG 1008, p. 28.)

For the reasons provided above and for those provided with respect to claim 1, Wexler in view of HTTP1.0 would have rendered claim 34 obvious.

c) Dependent Claims 2-6, 12, 14-18, and 35-42

Claim 2 depends from claim 1 and recites "wherein said intermediary device is a proxy server." A POSA would have understood that an intermediary device may exist between any two networked components in Wexler, such as the user's computer and the third party service. (GOOG 1003, ¶ 171.) Indeed, Wexler recognizes that a network such as the Internet "is comprised of many computers linked over telecommunication lines," and that a user may connect to the Internet through an Internet Service Provider. (GOOG 1007, 3:12-14 and 3:46-47.)

While Wexler does not explicitly disclose proxy servers as the intermediary devices, proxy servers were well known intermediary devices in systems using HTTP protocols. (GOOG 1003, ¶ 171.) As defined by HTTP1.0, a proxy is "[a]n intermediary program which acts as both a server and a client for the purpose of making requests on behalf of other clients." (GOOG 1008, p. 5.) It would have

been obvious for a POSA to use a proxy server as disclosed in HTTP1.0 as an intermediary between the user computer and the servers. (GOOG 1003, ¶ 171.) This would simply be a combination of prior art elements (the web-based information delivery system of Wexler and the proxy server of HTTP1.0) according to known methods to yield predictable results (a web-based information delivery system where communication passes through a proxy server). (*Id.*)

Claim 3 depends from claim 2 and recites "wherein said second portion of information is served from the primary or secondary servers." Claim 6 depends from claim 1 and recites, "wherein said second portion of information is served from one of the secondary servers." In Wexler, the primary server is the "third party server 13," the secondary server is the "advertiser's Web site 17," and the second portion of information is the "advertiser's Web page." (GOOG 1003, ¶ 174.) Thus, in Wexler, the second portion of information is served from the secondary server. (*Id.*)

Claim 4 depends from claim 1 and recites "wherein said first request signal is a content general request signal." Claim 35 depends from claim 34 and further recites, "wherein said banner request signal includes a content general Uniform Resource Locator address." It would have been obvious to a POSA that Wexler's download request signal 15a simply requests the advertiser's web page and the third party service 13 decides which particular web page of the advertiser to

display, thus making it a content general request signal. (*Id.*, ¶ 177.) As discussed with respect to claim 34, Wexler's banner request signal includes a URL. A POSA would have understood that, since the download request signal 15a in Wexler constitutes a content general request signal, the URL contained therein must also be content general. (*Id.*, ¶ 178.)

Claim 5 depends from claim 4 and recites "wherein said second request signal is a content specific request signal." Claim 40 depends from claim 35 and recites, "wherein said banner location signal includes a content specific Uniform Resource Locator address for the specified banner." These claims are rendered obvious over Wexler in view of HTTP1.0. (GOOG 1003, ¶¶ 181-182.) In Wexler, the redirect request 15b is a "redirect to the intended URL, *i.e.*, the advertiser's Web site." (GOOG 1007, 5:20-21.) As such, it identifies the specific URL at which the website is located. (GOOG 1003, ¶ 182.) It would have been obvious to a POSA that the download request signal 19a in Wexler is a content specific request signal that identifies the URL of the advertiser's web page. (*Id.*).

Claim 12 depends from claim 1 and recites "storing said second portion of information in the terminal." In Wexler, "the advertiser's Web page is downloaded to the user's Web browser 3." (GOOG 1007, 5:10-11.) A POSA would have understood that downloading information at the user's browser involves "storing." (GOOG 1003, ¶ 185.)

Claim 14 depends from claim 4 and recites "including selecting the composition of said second portion of information."[8] A POSA would have understood that third party service 13 in Wexler must first select which advertiser's web page to redirect the user's web browser to before the third party service 13 sends the redirect request 15b to the user's web browser. (GOOG 1003, ¶ 188.)

Claim 15 depends from claim 14 and recites "wherein the results of said composition selection are included in said location signal sent from the information server to the terminal."[9] Wexler's redirect request identifies the location of the advertiser's web page, which is the "result[] of said composition selection." (GOOG 1003, ¶ 191.)

Claim 16 depends from claim 1 and recites "wherein said location signal includes an HTTP 302 redirect command." Claim 39 depends from claim 34 and recites, "wherein said banner location signal constitutes an HTTP 302 redirect signal." Wexler explicitly states that redirect signal 15a is an HTTP 302 redirect signal. (GOOG 1007, 5:14-24.) HTTP1.0 also describes the use of a 302 signal. (GOOG 1008, p. 28.)

Claim 17 depends from claim 1 and recites "wherein said reference to a second portion of information includes at least a portion of a URL." Wexler's

---

[8] See fn. 6.

[9] See fn. 6.

hypertext file that embeds the banner 9 into the web page 7 contains a URL link to server 13. (GOOG 1003, ¶ 197.) Indeed, Wexler teaches that "[t]he aforementioned coding includes the URL pointing to the third party site 13." (GOOG 1007, 4:47-53.)

Claim 18 depends from claim 1 and recites "wherein said first portion of information is a web page, said second portion of information is a banner, and said reference is a link." As discussed with respect to claim 1, Wexler discloses the served web page 7 (with embedded banner 9) as the first portion of information. (GOOG 1003, ¶ 200.) The second portion of information is the advertiser's Web page that is served to the user's Web browser based on the initial banner 9. (*Id.*) Given the BRI, Wexler's advertiser's Web page falls within the definition of a "banner." (*Id.*) The banner 9 in web page 7 contains a link ("the reference") to the address of the third party accounting and statistical service 13. (*Id.*)

Claim 36 depends from claim 34 and recites "determining whether said specified banner is stored on the computer." Claim 37 depends from claim 36 and recites "wherein after said determining whether said specified banner is stored on the computer, if said specified banner is not stored on the computer then including causing a second banner request signal to be sent to said device requesting that said device serve said specified banner to the computer." Such limitations simply recite the default behavior of sending an HTTP request. (GOOG 1003, ¶¶ 203-204.)

HTTP1.0 states that any client or server may serve as a cache. (GOOG 1008, p. 6.) If a participant does not have a cached response, then any signal must be sent along to its destination. (GOOG 1003, ¶ 204.)

Claim 38 depends from claim 37 and recites "serving the specified banner from said device to said computer." As discussed with respect to claim 1, Wexler's downloading of the advertiser's web page to the user's web browser constitutes "serving the specified banner." (GOOG 1003, ¶ 208.)

Claim 41 depends from claim 34 and recites "tagging said specified banner as being cachable." Again, this is a standard functionality of HTTP. As described in HTTP1.0, "Expires entity-header field gives the date/time after which the entity should be considered stale. This allows information providers to suggest the volatility of the resource, or a date after which the information may no longer be val*Id.* Applications must not cache this entity beyond the date given." (GOOG 1008, p. 33.) That is, it would have been obvious to a POSA that a web browser application would cache a response (such as a banner) based on this tag if it is before the date given in the Expires entity-header field. (GOOG 1003, ¶ 211.) In fact, the '045 patent specification acknowledges that HTTP1.0's "Expiry tag" reads on this claim limitation. (GOOG 1001, 17:65-18:9.)

Claim 42 depends from claim 34 and recites "wherein said device is said primary server." Wexler, in its prior art section, discloses the scenario where the

first server and the second server are combined into a single server 17. (GOOG 1003, ¶¶ 214-215.) In such a system, "a banner typically points to the Web site 17 of the advertiser. A download request signal 19a is sent from the user's Web browser 3 to the advertiser's Web site 17. The Web site 17 downloads information, indicated by the reference numeral 19b, to the user's Web browser 3. The downloaded information is a copy of a hypertext source file operable to generate a Web page of the advertiser." (GOOG 1007, 4:10-18.) Accordingly, claims 2-6, 12, 14-18, and 35-42 would have been obvious over Wexler in view of HTTP1.0.

6. Ground 6: Claim 19 would have been obvious over Wexler in view of HTTP1.0 and further in view of Meeker

Claim 19 depends from claim 1 and recites "counting at least one display of said second portion of information on the terminal." The Wexler-HTTP1.0 combination discloses delivering advertising banners via web pages using the HTTP protocol, and counting advertisements on a per-click basis. (GOOG 1007, 2:57-59.) Meeker discloses mechanisms for the delivery and monitoring of Internet advertising content with the "Nuts and Bolts of Internet Advertising," that include the counting of both delivered banners (displays or impressions) and associated click-throughs. (GOOG 1010, pp. 6-2 and 6-5.) Because counting per impression and counting per click were the two standard methods to track Internet advertisements, it would have been obvious for a POSA to replace Wexler's per-

click counting with Meeker's per-impression counting, as such modification to Wexler would have simply been a design choice. (GOOG 1003, ¶ 220.)

Cost-per-impression was a standard advertising accounting metric at the time, and thus would have been a known option to a POSA. (*Id.*) Further, while both click-through and impression-based counting were known as described in Meeker, Meeker notes that the click-through model had not been adopted to the same degree as the impression model. (GOOG 1010, p. 6-5.) Accordingly, it would have been obvious to a POSA to use the impression/display model described in Meeker as a replacement for the click-based model of Wexler. (GOOG 1003, ¶ 220.) The replacement would also have been obvious to try because a POSA would only have to choose from a finite number of options (only two options: either counting clicks or counting impressions). (*Id.*) Accordingly, claim 19 would have been obvious over Wexler in view of HTTP1.0 and further in view of Meeker.

7. Ground 7: Claims 7-11 would have been obvious over Wexler in view of HTTP1.0 and further in view of Garland

Wexler in view of HTTP1.0 discloses delivering advertising banners through web pages and the HTTP protocol. The use of multiple servers, and the selection of a specific server to deliver information based on a "best-suited" criteria was well-known in the art. (GOOG 1003, ¶ 224.) For example, Garland discloses the concept of load balancing in an HTTP environment to "spread incoming request

load among several machines." (GOOG 1009, pp. 1 and 3.)

It would have been obvious to one of ordinary skill in the art to incorporate the load balancing of web resources as taught by Garland with the delivery of web based advertising banners using HTTP as disclosed by Wexler in view of HTTP1.0, so as to manage requests. (GOOG 1003, ¶ 225.) Such a combination would have been obvious for a skilled artisan as it is nothing more than a use of a known technique (optimizing information delivery over the Web as taught by Garland) for improving similar devices (the web-based advertising delivery system of Wexler/HTTP1.0) to yield predictable results (the advertisement provider computer selecting a least loaded (that is, best suited) advertisement provider computer to serve a banner). (*Id.*)

Claim 7 depends from claim 1 and recites "wherein after the primary server receives the first request signal from the terminal, further including determining which server connected to the computer network is best suited for serving said second portion of information to the terminal." As noted above, when Garland's distributed server system is incorporated into the advertisement delivery system of Wexler in view of HTTP1.0, it would have been obvious to determine which server is best suited for serving the advertiser's web page to the user. (*Id.*)

Claim 8 depends from claim 7 and recites "wherein results of said determining are included in said location signal sent from the information server to

the terminal." As discussed with respect to claim 15, Wexler's 302 redirect signal includes the specific location address of the advertiser's web site. (*See also*, GOOG 1003, ¶ 191.) Garland further teaches that the 302 redirect signal includes the address of the selected least loaded server. (GOOG 1003, ¶ 228; GOOG 1009, p. 4.) Because this specific location is only known after it has been selected using the "best-suited" criteria as described in Garland, the specific location is a "result[] of said determining" as recited in claim 8.

Claim 9 depends on claim 8 and recites "creating a matrix of selections between each of the terminals or groups of terminals and each of the servers and using said matrix to determine which of the servers is best suited to serve said second portion of information to the terminals or groups of terminals." As discussed with respect to claim 9 in Ground 2, Garland discloses this limitation. A POSA would have been motivated to use the measured results (including the measured latency) to select a server with the least latency according to a scheduling policy, as such scheduling policy further helps "reduce the latency of servicing a request." (GOOG 1009, p. 1; GOOG 1003, ¶¶ 231-232.)

Claim 10 depends from claim 9 and recites "wherein said selections contain round trip times between the servers and the terminals or groups of terminals." As discussed above, the measured latency in Garland includes the round trip time for each possible combination of servers and clients. (GOOG 1003, ¶ 235.)

Claim 11 depends from claim 1 and recites "including making one of the secondary servers a new primary server if the original primary server becomes inaccessible." It was well known in the art to use a group of web servers for handling failover situation such that functions of one web server are switched to a second server. (GOOG 1003, ¶ 238.) For example, Garland describes "**solv[ing] their problem of server failure**" through a round-robin resolution. (GOOG 1009, p. 8.) A POSA would have understood that using a group of web servers to solve server failure problem means that, if the first primary server fails, a second server can be made the new primary server. (GOOG 1003, ¶ 238.)

VIII. CONCLUSION

As set forth above, and as supported by the technical expert testimony of Peter Kent and Paul Leach, claims 1-12, 14-19, and 34-42 of U.S. Patent No. 6,286,045 are rendered obvious by the prior art cited herein. Petitioner has established a reasonable likelihood of prevailing on each ground, and prompt and favorable consideration of this Petition and institution of an *Inter Partes* Review are respectfully requested.

IX. MANDATORY NOTICES (37 C.F.R. § 42.8(a)(1))

The Real Party-In-Interest (37 C.F.R. § 42.8(b)(1)) is: Google Inc.

Petitioner Provides Notice of Related Matters (37 C.F.R. § 42.8(b)(2)):

Petitioner has concurrently filed two additional petitions for *inter partes*

review of the '045 patent, for different claims than those addressed in this petition. Petitioner has also concurrently filed two separate petitions for *inter partes* review of U.S. Patent No. 6,014,698, which is a continuation-in-part of the '045 patent. Petitioner is also involved in concurrent litigation in the U.S. District Court for the District of Delaware, case number 1:14-cv-00216.

Designation of Lead and Back-Up Counsel (37 C.F.R. § 42.8(b)(3)):

Lead Counsel	Back-Up Counsel
Michelle K. Holoubek (Reg. # 54,179) STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C. 1100 New York Avenue, NW Washington, DC 20005 202.772.8667 (telephone) 202.371.2540 (facsimile) mholoubek-PTAB@skgf.com	Michael V. Messinger (Reg. # 37,575) STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C. 1100 New York Avenue, NW Washington, DC 20005 202.772.8667 (telephone) 202.371.2540 (facsimile) mikem-PTAB@skgf.com

Notice of Service Information (37 C.F.R. § 42.8(b)(4)): Please direct all correspondence to lead counsel at the above address. Petitioner consents to email service at: MHoloubek-PTAB@skgf.com, and mikem-PTAB@skgf.com.

RESPECTFULLY SUBMITTED,

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.

Date: 2/2/15

Michelle K. Holoubek, Reg. No. 54,179
Attorney for Petitioner Google Inc.

1100 New York Avenue, N.W.
Washington, D.C. 20005-3934
(202) 371-2600

APPENDIX A - EXHIBIT LIST

***Google* Exhibit #**	***Description***
1001	Griffiths *et al.*, U.S. Patent No. 6,286,045 (filed May 19, 1997; issued September 4, 2001)
1002	File History for U.S. Patent No. 6,286,045
1003	Declaration of Peter Kent
1004	*Curriculum Vitae* of Peter Kent
1005	Declaration of Paul J. Leach
1006	*Curriculum Vitae* of Paul J. Leach
1007	Wexler, U.S. Patent No. 5,960,409 (filed October 11, 1996, issued September 28, 1999)
1008	Fielding *et al.*, "HTTP Working Group Internet Draft Hypertext Transfer Protocol – HTTP/1.0" (dated February 19, 1996; published February 20, 1996)
1009	Garland et al., "Implementing Distributed Server Groups for the World Wide Web" Carnegie Mellon University (January 25, 1995)
1010	Meeker, Mary, "Technology: Internet/New Media The Internet Advertising Report" Morgan Stanley, U.S. Investment Research (January 1997)
1012	Angles *et al.*, U.S. Patent No. 5,933,811 (filed August 20, 1996, issued August 3, 1999)
1013	Merriman *et al.*, U.S. Patent No. 5,948,061 (filed October 29, 1996, issued September 7, 1999)

Google Exhibit #	*Description*
1014	Davis *et al.*, U.S. Patent No. 5,796,952 (filed March 21, 1997, issued August 18, 1998)
1015	Returned Summons for Complaint *Richard A. Williamson, on behalf of and as Trustee for At Home Bondholders' Liquidating Trust* v. Google Inc., No. 1-14-cv-00216 (D.Del), filed February 19, 2014, summons served on February 20, 2014
1016	J. Mogul and P. Leach, "Simple Hit-Metering for HTTP" (dated January 21, 1997; published January 22, 1997)
1017	C. Brown and S. Benford, "Tracking WWW Users: Experience from the Design of HyperVisVR," Webnet 96, Oct. 15-19, 1996
1018	H. Skardal, "A Trip Report and some reflections," W3C Meeting on Web Efficiency and Robustness (Apr. 22, 1996)
1019	"WRQ Express PageMeter delivers breakthrough accuracy in Web metering," Business Wire, Apr. 28, 1997
1021	Lopez-Ortiz, *et al.*, "A Multicollaborative Push-Caching HTTP Protocol for the WWW" December 28, 1995
1022	Novak *et al.*, "New Metrics for New Media: Toward the Development of Web Measurement Standards" (September 26, 1996, published in the World Wide Web Journal, April 1997, 213-246)
1023	Pitkow *et al.*, "In Search of Reliable Usage Data on the WWW" (Proceedings of the Sixth International WWW Conference, April 1997)
1024	Mogul *et al.*, "HTTP Working Group Internet Draft Simple Hit-Metering and Usage-Limiting for HTTP,"(dated March 25, 1997; published March 27, 1997)

Google Exhibit #	*Description*
1025	Daniel W. Connolly, "Proposals for Gathering Consumer Demographics," (Created October 1995, updated November 6, 1995)
1026	Fielding *et al.*, "HTTP Working Group Internet Draft Hypertext Transfer Protocol – HTTP/1.1" (dated May 2, 1996; published May 3, 1996)
1027	Birrell *et al.*, "Grapevine: An Exercise in Distributed Computing" (April 1982)

CERTIFICATE OF SERVICE

The undersigned hereby certifies that on February 2, 2015, true and correct copies of the above-captioned **Petition for *Inter Partes* Review of U.S. Patent No. 6,286,045 Under 35 U.S.C. §§ 311-319 and 37 C.F.R. §§ 42.1-.80, 42.100-.123**, Exhibit List, and all associated exhibits were served in their entirety upon the following parties via FEDEX® Express:

FAEGRE BAKER DANIELS LLP
PATENT DOCKETING - INTELLECTUAL PROPERTY
2200 WELLS FARGO CENTER
90 SOUTH SEVENTH STREET
MINNEAPOLIS MN 55402-3901

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.

Date: 2/2/15

1100 New York Avenue, N.W.
Washington, D.C. 20005-3934

Michelle K. Holoubek, Reg. No. 54,179
Attorney for Petitioner Google Inc.

1946362_2.DOCX