IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent of Jiancheng Guo et al.	§	Attorney Docket No.: U022-0007RE
	§	
U.S. Patent 9,578,040	§	
	§	
Issue Date: February 21, 2017	§	Customer No.: 29,150
	§	
Filing Date: December 16, 2014	§	
	§	
For: PACKET RECEIVING METHOD, DEEP PACKET INSPECTION DEVICE AND SYSTEM	§	

Mail Stop "*Ex Parte* Reexam"
Attn: Central Reexamination Unit
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Commissioner:

DECLARATION OF ANGELOS KEROMYTIS, PH.D.

TABLE OF CONTENTS

I, Anglos Keromytis, declare as follows:

I. INTRODUCTION

1. I have been asked to submit this declaration on behalf of Unified Patents, LLC ("Requester") in connection with a request for *ex parte* reexamination of U.S. Patent 9,578,040 ("the '040 Patent," Ex. 1001) that issued from U.S. Application No. 15/572,514 ("the '514 Application"). Specifically, I have been retained as an independent expert consultant by Requester to provide my opinions on the technology claimed in, and the patentability or unpatentability of claims 1,4-6, and 9-11 of the '040 Patent ("the Challenged Claims"). Although I am being compensated at my usual rate of $550 per hour for the time I spend on this matter, no part of my compensation depends on the outcome of this proceeding, I have no financial interest in any of the parties, and I have no other interest in this proceeding.

II. QUALIFICATIONS

2. I have extensive experience in the area of computer and network security, and have been working in this field since 1993.

3. I am currently the John H. Weitnauer Technology Transfer Endowed Chair Professor with the School of Electrical and Computer Engineering (ECE) at the Georgia Institute of Technology (Georgia Tech). I am also a Georgia Research Alliance (GRA) Eminent Scholar, and an elected Fellow of the IEEE and of the ACM, the two premier professional organizations in the field of computing. I am also the co-founder and President of Voreas Laboratories Incorporated and of Aether Argus Incorporated, two Atlanta, GA-based technology startups in the area of cybersecurity, and a co-founder of Allure Security Technologies Inc., an MA-based technology startup also in the area of cybersecurity.

4. Before joining Georgia Tech, I was Program Manager with the Defense Advanced Research Projects Agency (DARPA), an R&D organization that is part of the Department of Defense, from 2014 to 2018. During that time, I conceived, initiated, and managed 5 major research programs (four of them in cybersecurity), and managed another 4, with a total budget of almost $500M. For my work at DARPA, I received the Department of Defense Superior Public Service Medal in 2018. Prior to DARPA, I served 1 year as Program Director with the National Science Foundation (NSF), responsible for the Secure and Trustworthy Cyberspace (SaTC) program, an $80M/year R&D effort that funds academic cybersecurity research across the country.

Prior to this tour for public service, I was an Associate Professor of Computer Science at Columbia University, as well as Director of the University's Network Security Laboratory. I joined Columbia in 2001 as an Assistant Professor, after receiving my M.Sc. and Ph.D. degrees in Computer Science, both from the University of Pennsylvania. My Ph.D. dissertation work was on the topic of secure access control for distributed systems and, in particular, on the management of trust in distributed computer networks.

5. I received my B.Sc. in Computer Science from the University of Crete, in Greece, in 1996. During my undergraduate studies, I worked as system administrator in the Computing Center at the University of Crete. Following that, I worked as network engineer at the first commercial Internet Service Provider ("ISP") in Greece, FORTHnet SA, where I was exposed to many network security issues.

6. I have actively participated in the Internet Engineering Task Force ("IETF"), a standards-setting body for the Internet, since 1995. In the late 1990s and early 2000s, my work with the IETF was primarily within the Internet Protocol Security ("IPsec") Working Group. In addition to contributing to the specification of the IPsec standards, I wrote the first implementation of the Photuris key management protocol (now RFC 2522). I also contributed to the first open-source implementation of the IKSAMP/IKE key management protocol for the open-source BSD operating system (now RFC 2409), and developed the first such implementation for the Linux operating system. My Linux implementation, named Pluto, was adopted by the National Institute of Standards and Technology ("NIST") in 1999. In addition, my implementation of IPsec for the open-source BSD operating system is currently used by many companies and governments around the world, and serves as the basis for several commercial products that employ cryptographic communications.

7. In 1999, I architected and implemented the first open-source framework for supporting hardware cryptographic accelerators. This framework is used in the open-source OpenBSD, NetBSD, FreeBSD, and Linux operating systems. My work in implementing firewalls and other cryptographic and network protocols has resulted in commercial systems and publications in refereed technical conferences and academic journals. I served as Working Group Secretary for the IETF IPsec Working Group (2003-2005) and as Security Area Advisor to the IETF at large (2003-2008).

8. In my position at Columbia University, I worked with a large group of graduate and postgraduate students in the area of cybersecurity. My past students now work in this field as university professors, as technical researchers for research laboratories, or as engineers for telecommunications companies. I have received federal, state, and corporate sponsorship to conduct cybersecurity research from the Department of Defense, the National Security Agency, the Defense Advanced Research Projects Agency ("DARPA"), the National Science Foundation, the Department of Homeland Security, the Air Force, the Office for Naval Research, the Army Research Office, the Department of the Interior, the National Reconnaissance Office, New York State, Google, Intel, Cisco, and others. In my 20 years as a professor, I have received over 54 million dollars to support my research in cybersecurity. I also regularly teach courses on cybersecurity, in addition to more general courses in computer science.

9. I have published over 250 technical papers in refereed journals, conferences, and workshops, all of which are directed to various areas of cybersecurity. I have also authored a book, coauthored another book, and contributed chapters for many other books that relate to cybersecurity. Between 1999 and 2010, I have drafted or co-drafted eight standards documents that were published as Request for Comments ("RFCs"). Several of these RFCs are directly related to IP security, like the RFCs I discuss in this declaration. For example, RFC 6042 relates to transport layer security; RFC 5708, RFC 2792, and RFC 2704 relate to key signature and encoding for trust management; and RFC 3586 relates to IP security policy requirements. Additionally, I am a coinventor on twelve issued U.S. patents, and have several other applications pending. Most of these patents and pending applications are related to network and systems security. I have chaired several international technical conferences and workshops in cybersecurity, including, for example, the International Conference on Financial Cryptography and Data Security (FC), ACM Computer and Communication Security (CCS), and the New Security Paradigms Workshop (NSPW). I have also served in over eighty technical program committees for such events. From 2004-2010, I served as Associate Editor for the premier technical journal on cybersecurity-the ACM Transactions on Information and Systems Security (TISSEC). Additionally, I have served on several advisory workshops to the United States Government on cybersecurity, including, among others, the Office of the Director of National Intelligence (ODNI)/National Security Agency (NSA) Invitational Workshop on Computational Cybersecurity in Compromised Environments (C3E) (2011), the Office of Naval Research (ONR) Workshop on Host Computer

Security (2010), the Intelligence Community Technical Exchange on Moving Target (2010), Lockheed Martin Future Security Threats Workshop (2009), and the ARO/FSTC Workshop on Insider Attack and Cyber Security.

10. My *curriculum vitae*, which is appended to the Request as Exhibit 1004, details my background and technical qualifications.

III. MATERIALS CONSIDERED

11. In forming my opinions expressed in this declaration, I have considered, among other things, the following documents. I understand the documents have been given the following exhibit numbers in this proceeding:

Exhibit	Description
Ex. 1001	U.S. Patent 9,587,040 to Guo et al. ("the '040 patent")
Ex. 1002	File History of the '040 Patent
Ex. 1005	Leech et al., "SOCKS Protocol Version 5," RFC 1928, March 1996 ("RFC 1928")
Ex. 1006	Koblas et al., "SOCKS," 1992, UNIX Security Symposium III Proceedings, 1992 ("Koblas")
Ex. 1007	Kitamura, "A SOCKS-based IPv6/IPv4 Gateway Mechanism," RFC 3089, April 2001 ("RFC 3089")
Ex. 1008	U.S. Patent Application Publication No. 2009/0157889 ("Treuhaft")
Ex. 1009	U.S. Patent Application Publication No. 2006/0167871 ("Sorenson")
Ex. 1010	U.S. Patent Application Publication No. 2004/0006621 ("Bellinson")
Ex. 1011	Subramanian et al., "An empirical vulnerability remediation model, IEEE International Conference on Wireless Communications, Networking and Information Security, 2010 ("Subramanian")
Ex. 1012	Excerpts from Microsoft Computer Dictionary, Fifth Edition, 2002
Ex. 1013	U.S. Patent 6,950,660 to Hsu et al. ("Hsu")
Ex. 1014	U.S. Patent Application Publication No. 2009/0049539 to Halbedel et al. ("Halbedel")

12. In forming my opinions, I have also relied on my education and experience.

IV. RELEVANT LEGAL STANDARDS

13. I am not an attorney. My analysis and opinions are based on my expertise in this technical field, as well as the instructions I have been given by counsel for the legal standards relating to patentability.

14. I have been informed by counsel for Requester that the following legal principles may apply to analysis of patentability based on 35 U.S.C. §§ 102 for anticipation and 103 for obviousness. I have also been informed that, in an *ex partes* reexamination proceeding such as this proceeding, a patent claim is unpatentable if it is shown by a preponderance of the evidence that the claim would have been anticipated by a prior art patent or publication, or obvious by one or more properly combined prior art patents or publications.

A. Level of Ordinary Skill in the Art

15. I have been instructed to consider patentability of the Challenged Claims through the lens of a person of ordinary skill in the art ("POSA") at the time of the claimed priority date of the'040 Patent--June 30, 2012. I am familiar with the level of ordinary skill in the subject matter of the '040 Patent in June 2012. Based on my review of the technology, and drawing on my own experience in the field, my analysis below assumes that a POSA would have had a bachelor's degree in computer science or computer engineering and two years of experience in computer and network security. In my opinion, however, less work experience may be compensated by a higher level of education, such as a master's degree, and vice versa.

16. Based on my qualifications discussed above, I qualified at least as a POSA of the '040 Patent by June 30, 2012.

17. My analysis below considers how a POSA would have understood the references listed above with respect to the Challenged Claims of the '040 Patent.

B. Anticipation

18. I have been informed a patent claim is unpatentable as anticipated under 35 U.S.C. § 102 if every limitation of the claimed invention is found in a single prior art reference--either expressly or required through inherency--as arranged in the claim.

C. Obviousness

19. I have been informed that, even if a single prior art reference does not disclose each and every element of a patent claim, the patent claim is still unpatentable as obvious under 35 U.S.C. § 103. It is my understanding that a claimed invention is unpatentable as obvious over a combination of prior art references if the differences between the claimed invention and the prior art are such that a POSA would have found the subject matter as a whole obvious.

20. I understand that obviousness is determined by evaluating: (1) the scope and content of the prior art, (2) the differences between the prior art and the claim, (3) the level of ordinary skill in the art, and (4) any secondary considerations of non-obviousness. To establish obviousness based on a combination of the elements disclosed in the prior art, it is my understanding that a challenger must provide a clear articulation of the reason(s) why the claimed invention would have been obvious. I understand this articulation may, but does not necessarily, require record evidence of an explicit teaching, suggestion, or motivation to combine the prior art in the way recited in a patent claim. Rather, prior art may be combined based on an express teaching, suggestion, or motivation from the prior art itself, or from a reasoned explanation of an expert witness or some other rationale.

21. For example, it is my understanding that this articulation can come from a number of rationales, which include but are not limited to (1) combining prior art elements according to known methods to yield predictable results; (2) simple substitution of one known element for another to obtain predictable results; (3) use of known technique to improve similar devices, methods, or products in the same way; (4) applying a known technique to a known device, method, or product ready for improvement to yield predictable results; (5) choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success, e.g., the combination is "obvious to try"; (6) known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art; and (7) some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed limitation.

22. I further understand that these rationales may be found explicitly or implicitly: (1) in the prior art; (2) in the knowledge of those of ordinary skill in the art that certain references, or disclosures in those references, are of special interest or importance in the field; or (3) from the nature of the problem to be solved. Additionally, I understand that the legal determination of the motivation to combine references allows recourse to logic, judgment, and common sense. In order to resist the temptation to read into prior art the teachings of the invention in issue, however, it should be apparent that the expert is not conflating "common sense" and what appears obvious in hindsight. I understand that if the teachings of a prior art would lead a POSA to make a modification that would render another prior art device inoperable, then such a modification may

not be obvious. I also understand that if a proposed modification would render the prior art invention being modified unsatisfactory for its intended purpose, then there may be no suggestion or motivation to make the proposed modification.

23. I understand that it may be improper to combine references where the references teach away from their combination. I understand that a reference may be said to teach away when a POSA, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the applicant. In general, a reference teaches away if it suggests that the line of development flowing from the reference's disclosure is unlikely to be productive of the result sought by the patentee. I understand that a reference teaches away, for example, if (1) the combination would produce a seemingly inoperative device, or (2) the references leave the impression that the product would not have the property sought by the patentee. I also understand, however, that a reference does not teach away if it merely expresses a general preference for an alternative invention but does not criticize, discredit, or otherwise discourage investigation into the invention claimed.

D. Claim Construction

24. Counsel has instructed me that, in an *ex parte* reexamination proceeding, the words of a claim are to be given their broadest reasonable interpretation consistent with the specification. Under this standard, I am instructed that the U.S. Patent and Trademark Office determines the scope of claims in patent applications not solely on the basis of the claim language, but upon giving claims their broadest reasonable construction in light of the specification as it would be interpreted by a POSA. It is my understanding that, because applicant has the opportunity to amend the claims during prosecution, giving a claim its broadest reasonable interpretation will reduce the possibility that the claim, once issued, will be interpreted more broadly than justified. In my analysis, I have applied this standard, as well as considered and applied any proposed constructions in the Request.

E. Substantial New Question of Patentability

25. Counsel has instructed me that the U.S. Patent and Trademark Office will order an *ex parte* reexamination proceeding if the prior art patents or printed publications submitted with the reexamination request raise a substantial new question of patentability with respect to the challenged claims. I am told that a prior art patent or printed publication raises a substantial question of patentability where there is a substantial likelihood that a reasonable examiner would consider the prior art patent or printed publication important in deciding whether or not the claim

is patentable. If the prior art would be considered important, then I understand the examiner should find a substantial new question of patentability so long as the same question of patentability has not already been decided as to the claim in a final holding by the Office or a federal court in an earlier review, or has not already been raised in another reexamination or other Office proceeding.

V. U.S. PATENT 9,578,040

A. Brief Overview of Deep Packet Inspection

26. The '040 Patent relates to a deep packet inspection (DPI) device. DPI is a type of data processing that inspects packets being sent over a computer network, and may take actions such as alerting, blocking, re-routing, or logging it accordingly. For example, in determining whether to take action on a packet, a DPI device may examine the headers (e.g., IP and/or TCP or UDP headers) and/or the data content of the packet.

B. Brief Overview of the Domain Name System

27. Certain aspects and functions of the DPI device of the '040 Patent, and the corresponding components of the prior art references discussed below, touch on the Domain Name System (DNS). DNS is the hierarchical naming system in which computers, services, or other resources connected to the Internet have an address represented as both a human-readable domain name (e.g., www.google.com) and a machine-readable Internet Protocol (IP) address (e.g., 111.222.333.444). DNS has been a basic, essential component of the functionality of the Internet since the late 1980s.

28. At a high level, the DNS includes a collection of name servers that maintain the domain name hierarchy and provide translation services between the domain name and IP address spaces. The name servers store address records mapping domain names to their corresponding IP addresses. A client application—such as a web browser—needs an IP address to access a resource connected to the Internet—such as a server hosting a website. Sometimes, however, the application does not have the particular IP address of the resource to which access is sought. Instead, the application may only have the corresponding domain name of that resource, as when a user enters the domain name (e.g., www.google.com) into a web browser.

29. To obtain the IP address of the Internet resource, the client application sends a DNS query to a name server in the DNS. The DNS query includes the domain name of the resource to which the application seeks access and requests the name server to provide the corresponding IP address. Upon receiving the DNS query, the name server "resolves," or converts, the domain into

its corresponding IP address by looking up the address record for the domain name in the DNS query, obtaining the corresponding IP address from the address record, and returning the IP address to the client application in a DNS response. As this point, the client application can access the Internet resource directly using the "resolved" IP address. This process of DNS resolution typically occurs transparently to the user of the client.

C. Summary of the '040 Patent

30. The '040 Patent is directed to a "a packet receiving method, a deep packet inspection and system." '040 Patent (Ex. 1001), 1:14-16. The deep packet inspection (DPI) device receives a service request packet sent by a terminal device. *Id.*, 3:27-28, FIG. 1 (step S101). The service request packet contains two pieces of information: (1) a terminal domain name indicating the terminal device; and (2) a server domain name indicating a server requested by the service request. *Id.*, 3:28-31. Upon receiving this service request, the DPI device determines whether to discard the service request packet or establish the requested connection. *See id.*, 3:66-5:46.

31. To do this, the DPI device first resolves (i.e., converts) the server domain name into its corresponding IP address, e.g., using DNS. *Id.* 3:66-4:12, FIG. 1 (step S102). Having resolved the IP address of the server to which the terminal device wants to connect, the DPI device now determines whether that IP address is contained in a preset list of allowable addresses for the terminal domain name. *Id.*, 4:13-5:47, FIG. 3 (step S103). If the IP address is not on the list, the DPI device discards (or denies) the service request, preventing the terminal device from accessing the IP address of the server. *Id.* But if the IP address is on the list, the DPI device establishes the requested connection. *Id.*

32. For the sake of reference, the claims for which reexamination is requested are reproduced below. I understand that claims 1, 6, and 11 are independent claims, while the remaining challenged claims depend directly or indirectly from these claims.

> 1. A packet receiving method, comprising:
> receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request packet sent by the terminal device;
> resolving the received server domain name to obtain a service server Internet protocol (IP) address; and
> discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list

the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device.

4. The method according to claim 1, wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:

if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device.

5. The method according to claim 1, wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:

if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device.

6. A deep packet inspection (DPI) device comprising a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method comprising:

receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request packet sent by the terminal device;

resolving the server domain name to obtain a service server Internet protocol (IP) address; and

discarding the packet if the service server IP address resolved does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with accessible service server IP addresses under an access authority of the terminal device.

9. The DPI device according to claim 6, wherein after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:

if the service server IP address resolved belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device.

10. The DPI device according to claim 6, wherein after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:

if the service server IP address resolved belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device.

11. A system, comprising:

a deep packet inspection (DPI) device; and

a terminal device, configured to send a service request packet to the DPI device, wherein the packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request sent by the terminal device;

the DPI device having a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method comprising:

receiving the service request packet sent by the terminal device;

resolving the server domain name received to obtain a service server Internet protocol (IP) address; and

discarding the packet if the service server IP address resolved does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with accessible service server IP addresses under an access authority of the terminal device.

33. In general, the claims of the '040 Patent relate to how the DPI device determines whether to discard a service request packet received from a terminal device or to establish the connection based on two pieces of information contained in the request—a domain name of the terminal and a domain name of the server. Figure 6 of the '040 patent, reproduced above, shows the system including the DPI device 30 and one or more terminal devices 40.

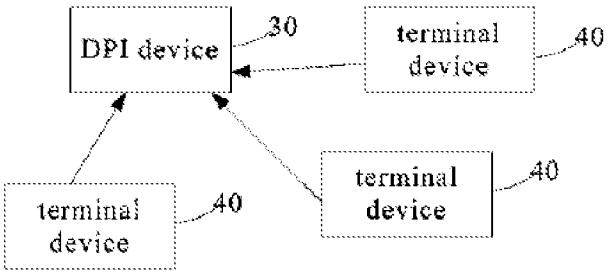

FIG. 6

34. Referring to the method shown in Figure 1 of the '040 Patent, reproduced below, when a terminal device wants to connect to a server, it sends a service request packet to the DPI device. *Id.*, 3:27-28, 6:34-25, FIG. 3 (step S101); *see also id.*, 6:34-38, FIG. 4 (step S204). The service request packet contains two pieces of information: (1) a terminal domain name indicating

the terminal device; and (2) a server domain name indicating a server requested by the service request. *Id.*, 3:28-31; *see also id.*, 6:34-38. The terminal domain name is simply "a unique identifier" the terminal device uses to identify itself to the DPI device from among "tens of thousands of terminal devices and service servers in the network." *Id.*, 3:32-51. And, similarly, the server domain name is simply a unique identifier of the server that the terminal device wishes to connect with. *Id.*, 3:51-65.

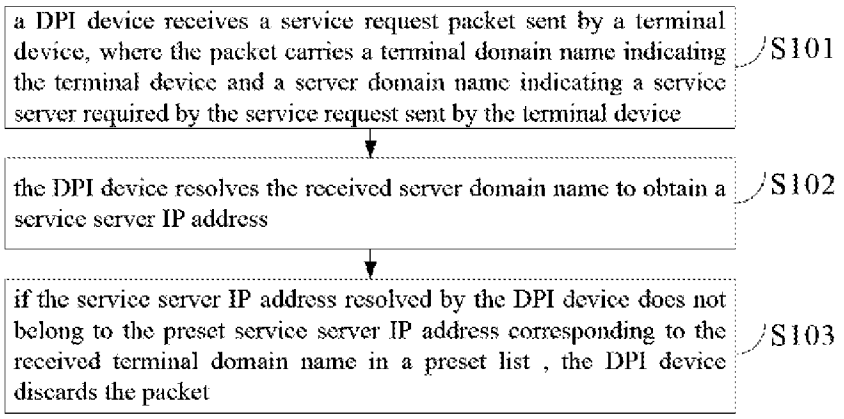

FIG. 1

35. Upon receiving the service request, the DPI device determines whether to discard it or establish the requested connection based on the two pieces of information in the request—the terminal domain name and server domain name. *See id.*, 3:66-5:4, 6:39-50. It is a two-step process:

36. First, the DPI device resolves (i.e., converts) the server domain name into its corresponding IP address. *Id.*, 3:66-4:12, FIG. 1 (step S102); *see also id.*, 6:39-41 (FIG. 3, step S205). For example, using DNS, the DPI device may convert the server domain name (e.g., www.google.com) to its corresponding machine-readable IP address (e.g., 2.2.2.2). *Id.*, 4:1-12.

37. Second, having resolved the IP address of the server to which the terminal device wants to connect, the DPI device now determines whether that IP address is contained in a preset list of accessible (i.e., allowable) addresses for the terminal domain name. *Id.*, 4:13-5:47, FIG. 3 (step S103); *see also id.*, 6:42-49, FIG. 3 (step S206). The "preset list is preset in the DPI device in advance" as configuration information before the method in Figure 1 occurs.[1] The '040 Patent gives an example of the preset list in Table 1 in the specification:

[1] In steps S201-203, the '040 Patent describes a process by which the DPI device establishes the preset list before the process in Figure 1, and steps S204-S206 of Figure 3, occur. *Id.*, 5:54-6:30, FIG. 3 (steps S201-S203). But establishing the preset list relates to unchallenged claims 2, 3, 7, and 8, so it is peripheral to the Request and I do not discuss it in detail here.

TABLE 1

Terminal domain name	Preset service server IP address
www.huawei.com	1.1.1.1 2.2.2.20
www.google.com	2.2.2.2

terminal domain name — corresponding accessible server IP address

Id., 4:27-34[2]. The left column lists terminal domain names of various terminal devices, and the right column lists the corresponding preset server IP addresses that the terminal domain names have authorization to access. *Id.* 4:35-5:24. For example, as highlighted above, a terminal device at the domain name www.google.com has authorization to access the corresponding server IP address 2.2.2.2 but not IP address 2.2.2.20. *Id.* A terminal device located at the terminal domain name www.huawei.com, on the other hand, may access the corresponding server IP address 2.2.2.20 but not 2.2.2.2. *Id.*

38. Continuing with step S103 of Figure 1, to determine whether to discard the IP address resolved from the server domain name contained in the request, the DPI device checks whether the preset list identifies the resolved IP address as an accessible IP address for the terminal domain name contained in the request.[3] Ex. 1001, 4:13-5:24; *see also id.*, 6:42-57, FIG. 3 (step S206). In Table 1 above, for example, if the request came from the terminal domain name www.google.com, the DPI device would check whether the resolved IP address is listed in the

[2] I designate annotated figures with "*"

[3] Perhaps stemming from translation of its original Chinese text to English, the '040 patent uses somewhat strange language to describe the scenarios when the preset list does or does not list the resolved IP address as an accessible IP address for a terminal domain name. Rather than state that the IP address is or is not a preset IP address on the list, the '040 patent respectively states that the resolved IP "belongs" or "does not belong" to a preset IP address on the list. *See, e.g.*, Ex. 1001, Abstract, 2:1-4, 2:15-19, 2:36-38, 4:16-17, 5:5-9, 5:32-35, 6:42-53, 6:63-66, 7:64-67, 8:34-38, 8:47-52, 8:55-59, 9:35-38, 10:12-15. In my opinion, A POSA have understood that, by "does not belong" or "belongs" to a preset address on the list, the '040 patent respectively means that the IP address is or is not listed as a preset address on the list.

right column as a corresponding accessible IP address for the terminal domain name www.google.com. *Id.*

39. If the resolved IP address is not on the preset list, the DPI device discards the request, preventing the terminal device from accessing the server at the resolved IP address. *Id.*, 4:13-5:17, FIG. 1 (step S103); *see also id.*, 6:50-57, FIG. 3 (step S206). For example, if the terminal domain name is www.google.com and the resolved IP address is 2.2.2.20, which is listed an authorized IP address for the terminal domain www.huawei.com but not for www.google.com, the DPI device would discard the request. *Id.*, 4:36-5:24. But if the resolved IP address is on the list for the terminal domain name—as the IP address 2.2.2.2 for the terminal domain name www.google.com---the terminal device has authorization to connect to the server and thus the DPI device establishes the connection. *Id.*, 6:36-7:5, FIG. 3 (step S207); *see also id.*, 5:17-24.

40. According to the '040 Patent, its technique of checking whether the resolved IP address is on the preset list for the terminal domain name, before establishing the connection, helps prevent a terminal device from fraudulently connecting to server by later changing its domain name when it connects to the server. *See id.*, 6:51-57, 8:39-46. Specifically, the '040 patent identifies a purported problem in which a terminal device, after obtaining the IP address of a server it is not authorized to access, changes its domain name in the host field of a connection request it later sends to IP address. *Id.* According to the '040 Patent, servers typically do not check the host field of a connection request to make sure the terminal device is authorized, and thus cannot prevent an unauthorized connection if the terminal device changes its domain name after it has already obtained the server's IP address. *See id.*, 1:26-48, 4:58-67, 5:10-17.

41. The '040 Patent gives an example in which a terminal device with the domain name www.google.com gains free access to a charged website—which the terminal domain name www.huawei.com has authorization to access but the terminal domain name www.google.com does not—by changing its domain name from www.google.com to www.huawei.com after obtaining the IP address of the server hosting the charged website. *See id.*, 1:26-48, 4:35-67, 5:5-17, 8:39-46. As shown in Figure 2, after obtaining the resolved IP address of the charged website, the terminal device alters its domain name from www.google.com to www.huawei.com in the host field of an HTTP GET request it sends to the resolved IP address. *Id.* 4:43-67, FIG. 2. According to the '040 Patent, by checking that the preset list contains the IP address of the server as authorized

for the terminal domain name, the DPI device can discard the request before the terminal device obtains the IP address of the server and attempts an unauthorized connection. *See, e.g.*, 6:51-57, 8:39-46.

D. Prosecution history

42. Although discussed in the Request, I provide a brief summary of the prosecution history (Exhibit 1002) of the '040 Patent for context. I understand that the examiner issued one office action before allowing the application, rejecting the claims as obvious over U.S. Patent 6,950,660 to Hsu et al. ("Hsu," Ex. 1013) in view of U.S. Patent Application Publication No. 2009/0049539 to Halbedel et al. ("Halbedel," Ex. 1014). Ex. 1002, 92-103.

43. As to the independent claims, I understand the examiner found that Hsu did not disclose "discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list," as recited in the independent claims, but that Halbedel did. *Id.*, 98. The examiner cited paragraph [0015] of Halbedel as disclosing this element. *Id.* In that passage, Halbedel explains that a hub maintains an access control list storing (1) a set of approved usernames and passwords authorized to access a particular server, or alternatively, (2) a set of valid IP addresses from which the server may be accessed. Ex. 1014, ¶ [0015].

44. In response to the rejections, I understand the applicant amended the independent claims to add to the "discarding" step the final "wherein" clause regarding the preset list. *Id.*, 1002, 42-47. For reference, the amendment to independent claim 1 is reproduced below:

> 1. (Currently Amended) A packet receiving method, comprising:
> receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request packet sent by the terminal device;
> resolving the received server domain name to obtain a service server Internet protocol (IP) address; and
> discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device.

Id., 43. Additionally, in the remarks, the applicant argued that the claimed preset list differed from

Halbedel's control list because it stored different information:

> ***Halbedel*** merely discloses a control list *that stores usernames, passwords, and user's IP addresses, in order to determine whether to grant the user access [to] the particular application server*.
> By contrast, as defined in amended claim 1, the preset list provides *the terminal domain name of each terminal device and a plurality of corresponding accessible service server IP addresses under an access authority of the terminal device*, so that after receiving a service request packet from the terminal device, resolving the domain name of a server carried in the service request packet, and obtaining IP addresses of the server, it can be determined whether the server is under access authority of the terminal device, by determining whether the IP address of the server resolved is in the preset list corresponding to the terminal's domain name.

Id. (emphasis in original)[4]. Following these amendments and arguments, I understand the examiner allowed the application, citing the "discarding" step with the added "wherein" clause in the examiner's statement of reasons for allowance. *Id.*, 24-31.

45. As I explain below, the prior art demonstrates that using a preset list, containing the terminal domain name of each terminal device and a plurality of corresponding authorized service server IP addresses, to control terminal device access to server IP addresses was well known years before the time of the '040 Patent. For example, as explained below, in the SOCKS protocol (Exhibits 1005-1007) developed in the mid-1990s, a SOCKS proxy server used a Configuration List mapping SOCKS client terminal domain names to corresponding accessible service IP addresses in determining whether to allow or deny incoming connections requests from terminal devices. Additionally, Treuhaft (Exhibit 1008) describes a system in which a DNS name server maintains subscriber information for each user or subscriber of the system identifying corresponding authorized/unauthorized server IP addresses that the user or subscriber is authorized/unauthorized to access. Upon receiving DNS queries from those users or subscribers, the DNS name server responds to the DNS queries accordingly based on the corresponding users' or subscribers' subscriber information.

VI. CLAIM CONSTRUCTION

A. "discarding the service request packet" (independent claims 1, 6, and 11)

46. Under the broadest reasonable interpretation standard, in my opinion, a POSA would have understood this language to include preventing unauthorized access to the resolved

[4] In this Declaration, emphasis is added unless otherwise specified.

service server IP address. The claim language itself supports this understanding by explaining that the service request packet received from the terminal device—and containing the server domain name from which the IP address was resolved—is discarded if the resolved IP address does not belong to (i.e., is not) a preset address in the preset list. *See* Ex. 1001, 10:43-48. In other words, because the preset list does not list the resolved IP address referenced in the service request packet as a preset (e.g., authorized) IP address, the service request packet is discarded and the requested connection is not granted, preventing the terminal device from accessing the resolved IP address.

47. Additionally, dependent claims 4 and 9 recite establishing the connection if the resolved service server IP address belongs to (i.e., is) a preset address on the list. Conversely, this suggests that the connection is not established if the resolved IP address is not a preset address on the list (i.e., the DPI device presents access). It follows that, preventing access when the resolved IP address is not on the list falls within the scope of independent claim 1, from which claims 4 and 9 depend. The specification of the '040 Patent also supports this understanding, explaining that if the resolved IP address is not on the preset list, "then the packet is considered to be abnormal, and the abnormal packet is discarded so as to prevent the terminal device A from successfully accessing the charged service through altering the packet without authorization[.]" Ex. 1001, 5:10-17. In other words, the service request is determined to be abnormal and discarded, preventing the terminal device from accessing the IP address, if the resolved IP address is not on the list.

B. "*if the resolved service server IP address does not belong to a preset service server IP address … in a preset list*" (independent claims 1, 6, and 11)

"*if the resolved service server IP address belongs to the preset service server IP address … in the preset list*" (dependent claims 4 and 5)

48. To the extent these phrases can be understood, I believe a POSA would have understood them to mean "if the resolved service server IP address is not a preset service server IP address … in a preset list" and "if the resolved service server IP address is a preset service server IP address … in the preset list," respectively. Perhaps stemming from translation of its original Chinese text to English, the '040 patent uses odd language to describe the scenarios in which the preset list does not or does not list the resolved IP address as a preset IP address for a terminal domain name. Rather than state that the IP address is or is not a preset address on the list, the '040 patent respectively states that the resolved IP "belongs" or "does not belong" to a preset address on the list. *See, e.g.*, Ex. 1001, Abstract, 2:1-4, 2:15-19, 2:36-38, 4:16-17, 5:5-9, 5:32-35,

6:42-53, 6:63-66, 7:64-67, 8:34-38, 8:47-52, 8:55-59, 9:35-38, 10:12-15. From the context, however, it is my opinion that a POSA would have understood that the '040 patent means the resolved IP address is not a preset address on the preset list when stating the resolved IP address "does not belong to a preset service server IP address … in a preset list," as recited in independent claims 1, 6, and 11. And, by the same token, I believe a POSA would have understood the '040 Patent means the resolved IP address is a preset address on the preset list when stating the resolved IP address "belongs to the preset server IP address … in the preset list."

VII. PRIOR ART PATENTS AND PRINTED PUBLICATIONS

49. In this Declaration I rely on the following references. I have been instructed by counsel to assume that each of these references legally qualifies as prior art against the '040 Patent.

Exhibit 1005, Leech et al., "SOCKS Protocol Version 5," RFC 1928, March 1996 ("RFC 1928")

50. RFC 1928 was published in 1996 in the RFC (Request for Comment) series. Ex. 1005, 1. Originally established in 1968, the RFC series edits and publishes technical and organizational documents about the Internet, including the specifications and policy documents produced by the Internet Engineering Task Force (IETF), the Internet Research Task Force (IRTF), the Internet Architecture Board (IAB), and Independent Submissions. As I explain above in the section regarding my background and qualifications, I have personal experience with RFC series because I have authored several published RFCs through my involvement in the IETF. Thus, I am quite familiar with the process of drafting a proposed RFC, submitting it to the RFC Editor for review, and seeing it published. Since long before the time of the '040 Patent, the RFC series constitutes a widely known and used resource for practitioners in the area of computer and network security (and, more broadly, Internet protocol standards). POSAs were aware of the RFC series and consulted its publications (such as RFC 1928 and RFC 3089) in the ordinary course of their work.

Exhibit 1006, Koblas et al., "SOCKS," 1992, UNIX Security Symposium III Proceedings, 1992 ("Koblas")

51. Published in 1992, Koblas is an article presented at the USENIX UNIX Security Symposium III conference, sponsored by the USENIX Association in cooperation with the Computer Emergency Response Team (CERT), on September 14-16, 1992 in Baltimore, MD.

Based on my experience in the field of computer and network security, POSAs are also familiar with the USENIX Association and CERT and would have attended this conference at which Koblas was presented.

Exhibit 1007, Kitamura, "A SOCKS-based IPv6/IPv4 Gateway Mechanism," RFC 3089, April 2001 ("RFC 3089")

52. RFC 3089 is another document in the RFC series, published in 2001.

Exhibit 1008, U.S. Patent Application Publication No. 2009/0157889 ("Treuhaft")

53. Treuhaft is a U.S. patent application in the area of network security filed in 2008 and published in 2009. Ex. 1008, 1.

Exhibit 1009, U.S. Patent Application Publication No. 2006/0167871 ("Sorenson")

54. Sorenson is a U.S. patent application in the area of network security filed in 2004 and published in 2006. Ex. 1009, 1.

Exhibit 1010, U.S. Patent Application Publication No. 2004/0006621 ("Bellinson")

55. Bellinson is a U.S. patent application in the area of network security filed in 2002 and published in 2004. Ex. 1010, 1.

Exhibit 1011, Subramanian et al., "An empirical vulnerability remediation model, IEEE International Conference on Wireless Communications, Networking and Information Security, 2010 ("Subramanian")

56. Subramanian is an article presented at the "2010 IEEE International Conference on Wireless Communications, Networking and Information Security," a conference held by Institute of Electrical and Electronics Engineers (IEEE) in June 2010 in Beijing, China. I am told that, following the conference, Subramanian was uploaded to the IEEE Xplore digital library on August 5, 2010. Long before the '040 Patent, POSAs were familiar with IEEE, attended IEEE conferences, and consulted IEEE publications in the ordinary course of their work. Thus, in my opinion, POSAs would have attended the conference at which Subramanian was presented. And, if interested in the subject matter or author, POSAs could have obtained Subramanian in Xplore.

57. I understand that none of the references I rely upon in this Declaration was cited or considered during the prosecution of the '040 Patent.

VIII. GROUND 1: RFC 1928 IN VIEW OF KOBLAS AND RFC 3089 (THE "SOCKS REFERENCES") PRESENTS A SUBSTANTIAL NEW QUESTION OF PATENTABILITY

58. RFC 1928, Koblas, and RFC 3089 all relate to the SOCKS protocol—an Internet protocol for exchanging network packets over TCP/IP between a client and server through a proxy server, called a SOCKS proxy server. *See* Exs. 1005, 1006, 1007. As they all refer to the same protocol, in my opinion, POSAs would have considered these together when considering whether and how to use a system like SOCKS. Accordingly, in this Declaration I sometimes refer to RFC 1928, Koblas, and RFC 3089 collectively as "the SOCKS references."

59. In my opinion, the SOCKS references raise a substantial new question of patentability as to claims 1, 4-6, and 9-11 because their a patent examiner would have considered their teachings to be important in deciding whether the Challenged Claims are patentable. For example, as discussed below here and in the Detailed Discussion section, it is my opinion that the SOCKS references, when considered as an ordered combination, teach each limitation of the claims, including the supposedly novel feature that led the examiner to allow the claims: "discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device."

A. Overview of the SOCKS References

1. RFC 1928

60. SOCKS is an Internet protocol that exchanges network packets over TCP/IP between a client and server through a proxy server, called a SOCKS proxy server. RFC 1928 describes SOCKS Protocol Version 5 and aims to "provide a general framework … to transparently and securely traverse a firewall" in SOCKS[5]. Ex. 1005, 1; *see also id.*, 2 ("The [SOCKS] protocol described here is designed to provide a framework for client-server applications

[5] In SOCKS, the SOCKS proxy server is sometimes called a "firewall". I refer to this SOCKS firewall component as the "SOCKS server" or "SOCKS proxy server," and use these terms interchangeably throughout the Declaration.

in both the TCP and UDP domains to conveniently and securely use the services of a network firewall.").

61. When a SOCKS client wishes to connect to a server or other object behind the SOCKS proxy server, it sends a connection request to the SOCKS proxy server:

> When a TCP-based client wishes to establish a connection to an object that is reachable only via a firewall (such determination is left up to the implementation), it must open a TCP connection to the appropriate SOCKS port on the SOCKS server system. The SOCKS service is conventionally located on TCP port 1080. If the connection request succeeds, the client enters a negotiation for the authentication method to be used, authenticates with the chosen method, then sends a relay request. The SOCKS server evaluates the request, and either establishes the appropriate connection or denies it.

Id., 2-3; *see also id.*, 4 ("Once the method-dependent subnegotiation has completed, the client sends the request details."). The SOCKS connection request has the following form:

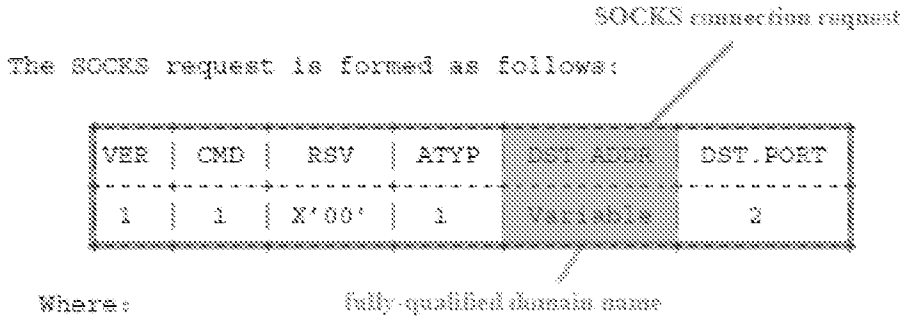

Id., 4*. As highlighted in the Figure above, the SOCKS request has a DST.ADDR field, which contains the "desired destination address" sought by the SOCKS client. *Id.* The preceding ATYP field specifies the "address type of [the] following address" contained in the DST.ADDR field, *id.*, 4, which can take the form of "a fully-qualified domain name," *id.*, 5. Additionally, the SOCKS requests contains a source IP address of the SOCKS client because the connection request is sent

over TCP/IP, which requires all packets to specify the source and destination IP addresses.

62. RFC 1928 teaches that "The SOCKS server evaluates the request, and either establishes the appropriate connection or denies it." *Id.*, 3. But because RFC 1928 focuses on laying out a general framework and protocol format for interacting with a SOCKS proxy server, RFC 1928 itself does not expressly describe the mechanism for evaluating connection requests.

2. Koblas

63. Koblas uses SOCKS to address "[o]ne of the more important [security] issues" when connecting to a network over the Internet: "intruders attempting to gain access to local hosts." Ex. 1006, 3[6]. Koblas proposes "several strategies which can be used to configure an Internet connection to prevent unwanted intrusion" using the SOCKS protocol. *Id.*, 3.

64. In one strategy, Koblas teaches that the SOCKS server uses a "Configuration File" to allow or deny the connection request. The Configuration File contains an entry for each SOCKS client source identifying corresponding "permit" or "deny" destination addresses to which the SOCKS server will respectively permit or deny connections requested by the SOCKS client source:

> The configuration file is located on the firewall host and is used by sockd when determining whether to accept or deny requests. The file is parsed from beginning to end, with the first fully matching line returning the accessibility. The syntax of the lines in this file is as follows:
>
> {permit | deny} <source-host> <mask> [<<dest-host> <mask> [<operator> <port>]]
>
> source address destination address
>
> Lines begin with either 'permit' or 'deny' following which are either 2, 4, or 6 fields, containing host address and mask pairs for source and destination, as well as a boolean operator and a service port.

Id., 7*.

[6] For ease of reference, I cite the PDF page number of Exhibit 1006.

65. As highlighted above, in the SOCKS Configuration File, the <source-host> field contains the host address of the SOCKS client source (claimed terminal device) and the corresponding <dest-host> field in the entry contains the address of the corresponding destination server (claimed corresponding preset service server IP address). *Id.* Depending on which value the {permit | deny} field contains, the SOCKS server will respectively permit or deny the SOCKS client source named in the <source-host> field to connect to the corresponding destination server address in the <dest-host> field. *Id.*, 1006, 7. Additionally, Koblas explains that "[h]ost addresses and services may be specified either by name or number," meaning SOCKS supports listing the addresses in the <source-host> and <dest-host> fields as a domain name or an IP address. *Id.* 1006, 8.

66. In Figure 5, Koblas "shows an example of how the lines in a configuration file might appear":

FIGURE 5. A Sample Configuration File

```
#
# Deny all host to every host whois service
#
deny 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq whois
#
# Let lloyd.mips.com only use finger service to sgi.com
#           source                destination
permit lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0 eq finger
deny lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0
#
# Allow all hosts on the 130.62 network access to the world
#
permit 130.62.0.0 0.0.255.255
#
# Deny all hosts which do not match anything in this file
# (i.e. All hosts coming in from the Internet)
#
```

Id., 8. In this Sample Configuration file, the SOCKS server permits a request from the source device with the domain name lloyd.mips.com to connect to the corresponding destination sgi.com, and denies requested connections "which do not match anything in this file." *Id.* Although shown as a domain name in this example, the server destination address sgi.com "may be specified either by name or number," so the IP address of sgi.com could be used instead. *Id.*, 8.

3. RFC 3089

67. RFC 3089 explains that, "[i]n all communication applications, it is [] necessary to obtain destination IP address information to start a communication." Ex. 1007, 4. To that end RFC 3089 describes a process by which the SOCKS server uses DNS to resolve the fully qualified domain name (FQDN) of the destination node (Destination D) into its "real IP address" when receiving a connection request from the source node (Client C):

> The detailed internal procedure of the "DNS name resolving delegation" and address mapping management related issues are described as follows.
>
> 1. An application on the source node (Client C) tries to get the IP address information of the destination node (Destination D) by calling the DNS name resolving function (e.g., gethostbyname()). At this time, the logical host name ("FQDN") information of the Destination D is passed to the application's *Socks Lib* as an argument of called APIs.
> 2. Since the *Socks Lib* has replaced such DNS name resolving APIs, the real DNS name resolving APIs is not called here. The argued "FQDN" information is merely registered into a mapping table in *Socks Lib*, and a "fake IP" address is selected as information that is replied to the application from a reserved special IP address space that is never used in real communications (e.g., 0.0.0.x). The address family type of the "fake IP" address must be suitable for requests called by the applications. Namely, it must belong to the same address family of the Client C, even if the address family of the Destination D is different from it. After the selected "fake IP" address is registered into the mapping table as a pair with the "FQDN", it is replied to the application.
> 3. The application receives the "fake IP" address, and prepares a "socket". The "fake IP" address information is used as an element of the "socket". The application calls socket APIs (e.g., connect()) to start a communication. The "socket" is used as an argument of the communication APIs (e.g., to send and receive data).
> 4. Since the *Socks Lib* has replaced such socket APIs, the real socket function is not called. The IP address information of the argued socket is checked. If the address belongs to the special address space for the fake address, the matched registered "FQDN" information of the "fake IP" address is obtained from the mapping table.
> 5. The "FQDN" information is transferred to the *Gateway* on the relay server (Gateway G) by using the SOCKS command that is matched to the called socket APIs. (e.g., for connect(), the CONNECT command is used.)
> 6. Finally, the real DNS name resolving API (e.g., getaddrinfo()) is called at the *Gateway*. At this time, the received "FQDN" information via the SOCKS protocol is used as an argument of the called APIs.
> 7. The *Gateway* obtains the "real IP" address from a DNS server, and creates a "socket". The "real IP" address information is used as an element of the "socket".
> 8. The *Gateway* calls socket APIs (e.g., connect()) to communicate with the Destination D. The "socket" is used as an argument of the APIs.

Id., 5-6.

B. The Combination of RFC 1928, Koblas, and RFC 3089 Presents a Substantial New Question of Patentability

68. In my opinion, the combination of the SOCKS references presents a substantial new question of patentability with respect to the '040 Patent. SOCKS, as described in RFC 1928, is similar to the claimed invention, in part because it uses a SOCKS proxy server in a firewall-like role similar to the DPI device in the '040 Patent. Like the DPI device in the '040 Patent, the SOCKS server receives a connection request from a source host device seeking access to a destination host server. The SOCKS connection request, like the service request in the '040 Patent, contains two pieces of information: (1) an identifier of the source device; and (2) a domain name of the destination server. Both the SOCKS server and the DPI device resolve the domain name of the destination server into its corresponding IP address. And, like the DPI device, RFC 1928 teaches that the SOCKS server evaluates the connection request and denies it if the request is not appropriate.

69. Although RFC 1928 does not expressly describe the mechanism to evaluate the connection request, Koblas discloses that the SOCKS server uses a Configuration File for this purpose. In my opinion, the SOCKS Configuration File goes directly to the purported point of novelty of the '040 Patent and the reason the examiner allowed the Challenged Claims. Namely, like the "preset list" of the '040 Patent, the Configuration File of Koblas lists the domain name of each source device and corresponding accessible server IP addresses under an access authority of the source device. When a connection request is received, the SOCKS server applies the Configuration File to determine whether to allow or deny the connection. If the Configuration File lists the IP address of the destination server as an allowed address for the domain name of the source host making the connection request, the SOCKS server allows the connection. If not, however, the SOCKS server denies the connection. As explained below in the Detailed Discussion section for the specific claim elements pertinent to the combination, a POSA would have been motivated to combine the SOCKS references and had a reasonable expectation in doing so.

Accordingly, as described above and below in the Detailed Application section, it is my opinion that the SOCKS references disclose, or at least render obvious, "*discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a*

plurality of accessible service server IP addresses under an access authority of the terminal device," as recited by independent claims 1, 6, and 11. *Id.* Thus, the SOCKS references present a substantial new question of patentability with respect to the Challenged Claims.

IX. GROUNDS 2-5: TREUHAFT, TREUHAFT IN VIEW OF SORENSON, AND TREUHAFT/SORENSON IN VIEW OF BELLINSON PRESENT SUBSTANTIAL NEW QUESTIONS OF PATENTABILITY

A. Overview of the References

70. As explained below, Treuhaft, Treuhaft in view of Sorenson, and Treuhaft/Sorenson in view of Bellinson present substantial new questions of patentability as to claims 1, 4-6, and 9-11 because a reasonable examiner would consider their teachings to be important in deciding whether or not the '040 Patent claims are patentable. For example, as discussed below here and in the Detailed Discussion section, Treuhaft, Treuhaft in view of Sorenson, and Treuhaft/Sorenson in view of Bellinson, when considered as an ordered combination, teach each limitation of the claims, including the purportedly novel feature that led the examiner to allow the claims: "discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device."

1. Treuhaft

71. Treuhaft describes a system in which a DNS name server 120 of Treuhaft maintains subscriber information 208 for various users or subscribers of the system. *See* Ex. 1008, ¶¶ [0028], [0029], [0034], [0036], [0039], [0054], [0060], [0064], FIG. 2 (subscriber information 280). "The subscriber information can include preferences or other settings for how a user or subscriber wishes to control domain name resolution within the DNS resolution features." *Id.*, ¶ [0060]. "For example, a user or subscriber may establish subscriber information that instructs DNS nameserver 120 to alter responses to DNS requests that are associated with adult web sites, potential phishing or pharming sites, and other sites deemed inappropriate by the user or containing material illegal in the country of the user." *Id.*, ¶ [0028]. Thus, the DNS name server 120 has a similar role to the SOCKS server and the DPI device of the '040 patent.

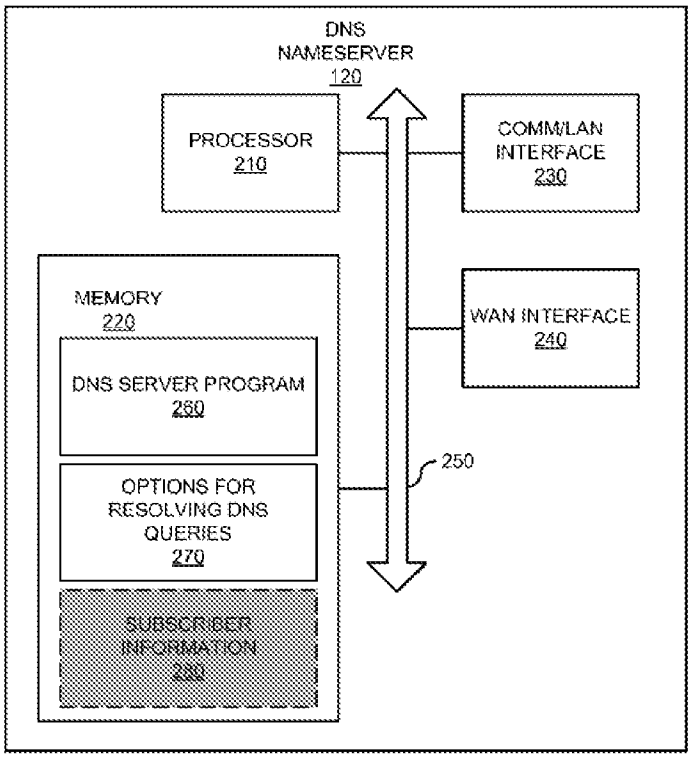

FIG. 2

Treuhaft, FIG. 2*

72. In Treuhaft, the DNS name server 120 receives a DNS query from a host device 105 seeking an IP address to connect to a server. *Id.*, ¶¶ [0063], [0064], FIG. 5A (step 525), FIG. 5B (step 530). The DNS query contains two pieces of information (1) a domain name for a URL of the server; and (2) "control information" identifying the host device 105:

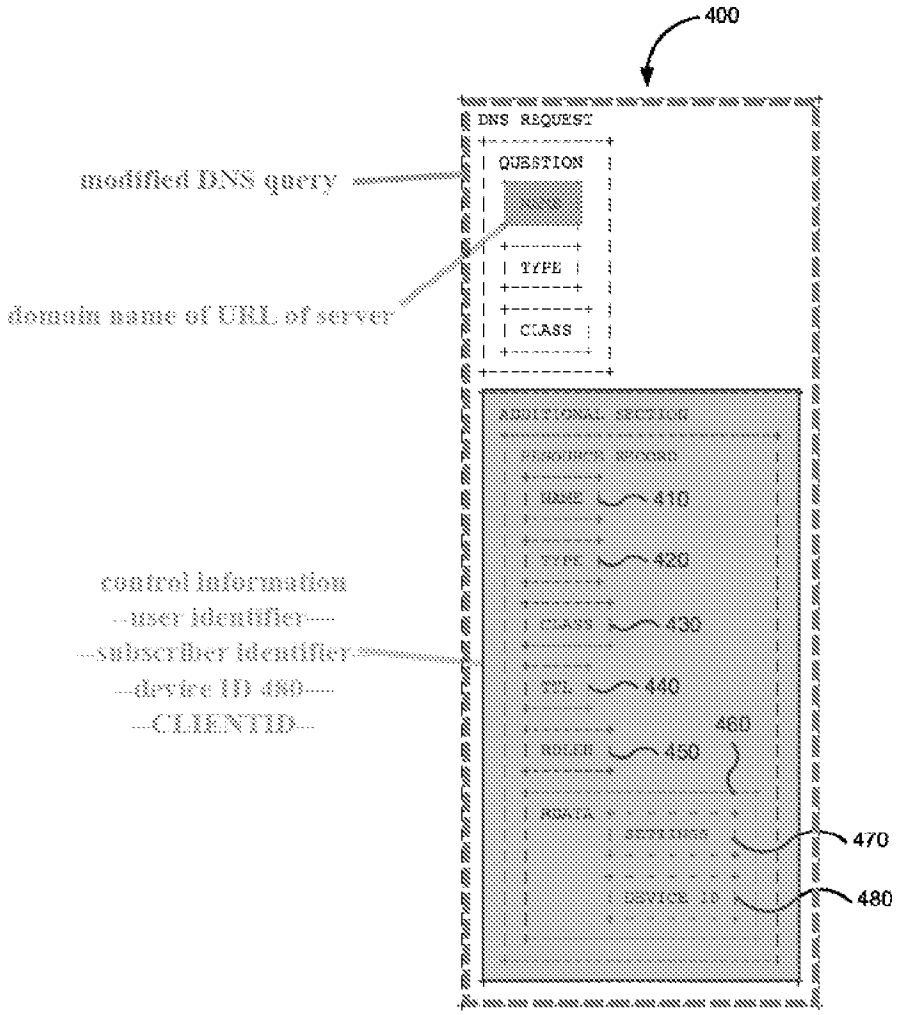

Treuhaft, FIG. 4*

73. As highlighted in Figure 4 above, the DNS query 400 includes a NAME field containing the domain name of the URL the host device 105 seeks to access. *Id.*, ¶ [0054]. Additionally, before the host device 105 sends the DNS query to name server 102, "the DNS query is modified with control information." *Id.*, ¶ [0058], FIG. 5A (step 520); *see also id.*, ¶ [0067] ("control information may be encoded into an individual DNS query that enables a DNS nameserver to identify DNS resolution options, filters, or features to apply when resolving the individual DNS query"). "The control information may specify … a user or subscriber identifier, a device identifier, or the like." *Id.*, ¶ [0036].

74. Upon receiving the DNS query, the "DNS nameserver 120 determines how to respond to host device 105" by applying the subscriber information 280 for the particular user or subscriber. *Id.*, ¶ [0032]; *see also id.*, ¶¶ [0028], [0029], [0034], [0036], [0039], [0054], [0060],

[0064]. Specifically, the DNS name server 120 parses the control information in the DNS query to identify the particular user or subscriber associated with the DNS query, and then retrieves that user or subscriber's subscriber information 280. *Id.*, ¶ [0064], FIG. 5B (step 353); *see also id.*, ¶ [0060].

75. Using the user or subscriber's subscription information 280, the DNS name server 120 "make[s] a decision whether to use the corresponding IP address or another IP address when generating a DNS response based on applying one or more DNS resolution options or features". *Id.* ¶ [0065]. For example, rather than return the resolved IP address requested by the host device 105, "DNS nameserver 120 may determine to substitute the IP address of a website that provides information why the domain name is being block[ed], forwarded, filtered, or otherwise includes material the user has expressed a desire to control." Ex. 1008, ¶ [0065], FIG. 5B (steps 540). Then, the DNS name server 120 generates a DNS response "substitut[ing] [the] IP address based on applying one or more of the available DNS resolution options, filters, or features" and sends the DNS response with the substituted IP address to the host device 105. Ex. 1008, ¶ [0066], FIG. 5B (steps 545, 550).

2. Sorenson

76. Sorenson discloses a "system and method for blocking access by a network device to specific network resources by comparing a specific resource identifier against entries in a blacklist and facilitating a connection accordingly." Ex. 1009, Abstract. In Sorenson, the system receives "a call request for the establishment of a communication session between IP device 12 and associated service 20. *Id.*, ¶ [0027]; *see also id.*, ¶¶ [0031], [0032]. The call request "include[s] a specific identifier such as an entered IP address, domain name, or conventional phone number or name resolved into one of an IP address or domain name." *Id.*, ¶ [0031].

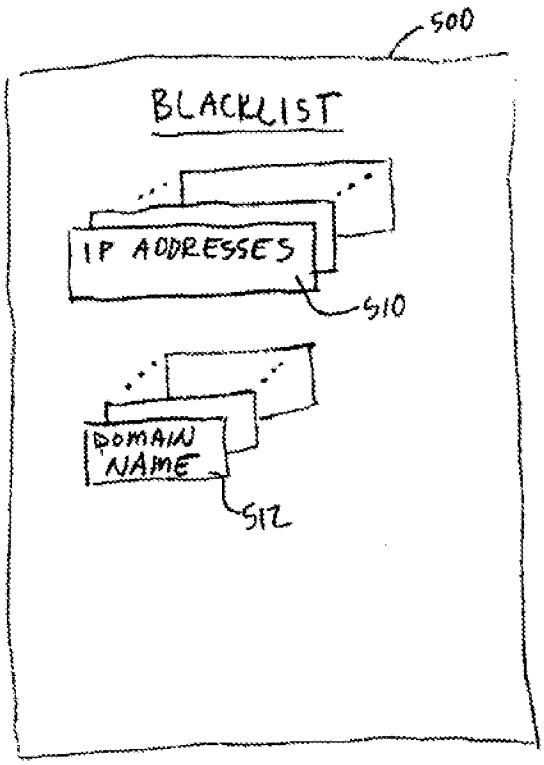

77. Before granting the call request, the system of Sorenson performs a two-stage blacklist check to determine whether to establish the

connection or discard the request. *See id.*, ¶¶ [0031]-[0032], FIG. 6. As shown in Figure 2 of Sorenson reproduced to above, the blacklist 500 contains both blacklisted domain name names 512 and blacklisted IP addresses 510. *Id.* 1009, ¶ [0028], FIG. 2.

78. First, as highlighted in step 610 of Figure 6 reproduced below, Sorenson performs a domain-name-blacklist check by "compar[ing] 610 the domain name against the blacklist 500´ (FIG. 2) to determine 612 if the domain name is located within the blacklist 500´." *Id.*, ¶ [0031]. "If the domain name utilized for initiating the call is located with the blacklist 500´, then the IP device denies 618 the completion of the call and may alternatively notify the user of such denial." *Id.* But "[i]f the domain name is not on the blacklist, then" Sorenson resolves the address and performs a second, IP-address-blacklist check before establishing the connection. *Id.*

79. Specifically, Sorenson "resolves … the domain name into an IP address for further comparison" in step 614, *id.*, and then "compares … the IP address against the blacklist 500'" in step 616, *id.*, ¶ [0032]. If the IP address is located within the blacklist 500', Sorenson denies the connection. *Id.*, ¶ [0032], FIG. 6 (step 618). If the IP address is not found on the blacklist 500', however, Sorenson establishes the connection. *Id.*, ¶ [0032], FIG. 6 (step 622).

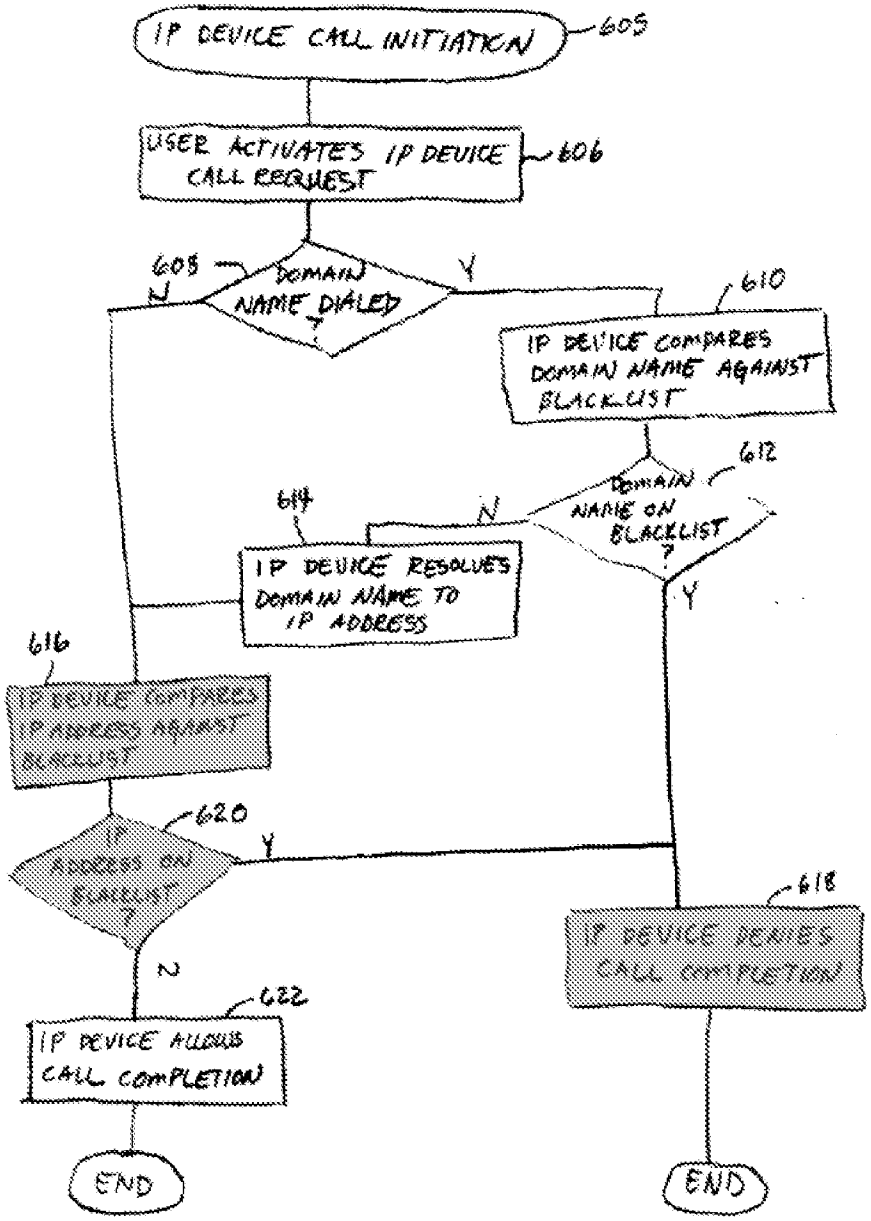

Sorenson, FIG. 6*

3. Bellinson

80. Similar to '040 Patent, SOCKS, Treuhaft, and Sorenson, Bellinson discloses "a system and method for controlling whether a user may access certain Internet sites" by applying "an allow-block list" to "determine[] whether the URL is referenced on the allow-block list and, if so, allow[] or disallow[] access to the site referenced by the URL accordingly." Ex. 1010, Abstract. In Bellinson, "[t]he allow-block list is a listing of specific site identifiers that the user is expressly authorized to view or prohibited from viewing." *Id.*, ¶ [0020]; *see also id.*, ¶¶ [0009]

("The allow-block list is a file containing a listing of specific URLs that the user is expressly authorized to view or expressly prohibited from viewing.").

81. As shown in Figure 3, reproduced to the right, in step 244 Bellinson's system receives an access request containing "a specified site identifier that references an Internet site. Examples of such site identifiers include designators such as www.microsoft.com but could also include an Internet Protocol (IP) address." *Id.*, ¶ [0049], FIG. 3. In step 246 of Figure 3, Bellinson "determines whether the site identifier is on the allow-block list at step 246. *Id.* "If the site identifier is referenced on the allow-block list," in step 248 Bellinson "determine[s] whether the site identifier is designated as blocked on the allow-block list." *Id.* "If the site identifier is [designated as] blocked," Bellinson blocks the connection. *Id.*, ¶ [0050]. Otherwise, Bellinson allows the connection. *Id.*, ¶ [0051].

FIG. 3.

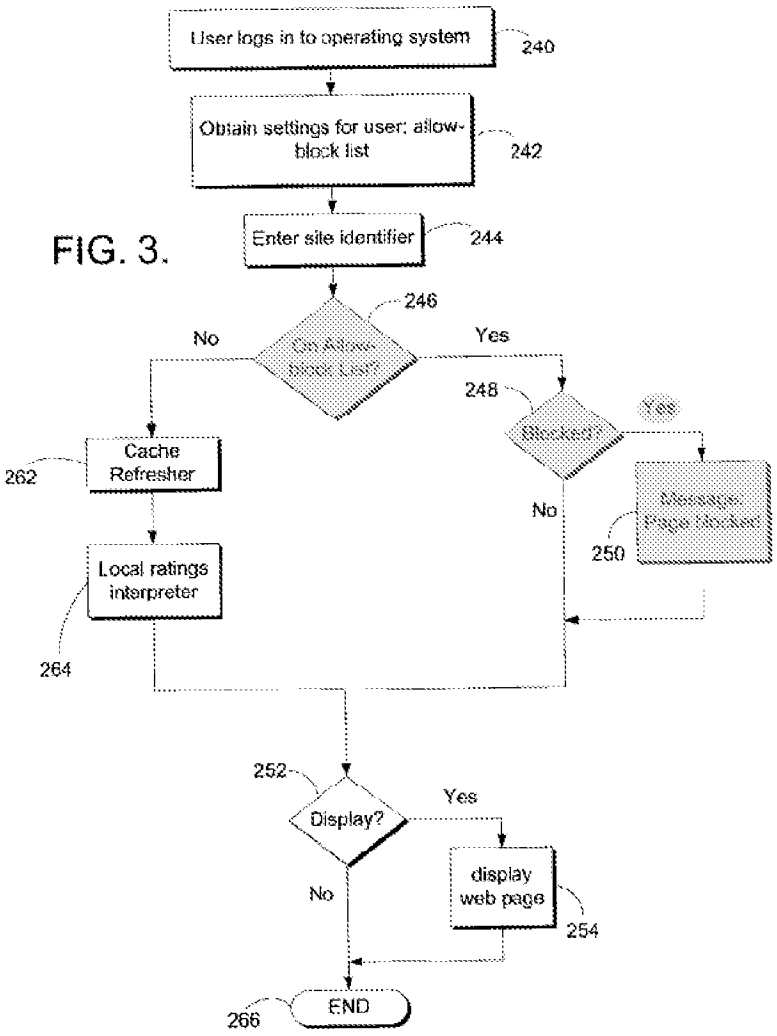

B. Treuhaft, Treuhaft in view of Sorenson, and Treuhaft/Sorenson in view of Bellinson Present Substantial New Questions of Patentability

82. In my opinion, combinations of Treuhaft, Treuhaft and Sorenson, and Treuhaft/Sorenson in view of Bellinson raise substantial new questions of patentability with respect to the Challenged Claims of the '040 Patent. Treuhaft is similar to the system claimed in the '040 Patent, in part, because it uses a DNS name server 120 in a role similar to the DPI device in the '040 Patent. Like the DPI device in the '040 Patent, Treuhaft's DNS name server 120 receives a request (a modified DNS query) from a host device seeking to access a destination

server. The modified DNS query, like the service request in the '040 Patent, contains two pieces of information: (1) control information identifying the host device; and (2) a domain name of the destination server the host device seeks to access. Both Treuhaft's DNS name server and the DPI device resolve the domain name of the destination server into its corresponding IP address. And, similar to the DPI device, Treuhaft's DNS name server evaluates the DNS query and denies it if the request is not appropriate.

83. In fact, in my opinion, the manner in which Treuhaft's DNS name server evaluates the DNS query goes directly to the purported point of novelty of the '040 Patent and the reason the examiner allowed the '040 Patent claims. Like the DPI devices applies the "preset list" in the '040 Patent, the DNS name server of Treuhaft applies subscriber information 280 to determine whether to allow or deny the requested connection. And Treuhaft discloses or suggests that the subscriber information 280, like the "preset list" of the '040 Patent, identifies corresponding server IP addresses under the access authority of each user or subscriber in Treuhaft. Accordingly, in my opinion, Treuhaft alone presents a substantial new question of patentability with respect to the '040 Patent.

84. In my opinion, Sorenson and Bellinson further address the purported point of novelty of the '040 Patent and raise substantial new questions of patentability in combination with Treuhaft. For example, Sorenson's blacklist of IP addresses and Bellinson's allow-block list of addresses each correspond to the claimed preset list. It is noted that element [1.4] recites that the preset list contains addresses "under an access authority of the terminal device. Though Sorenson's IP address blacklist identifies corresponding addresses not under the access authority of the device seeking the connection, as evidenced by Bellinson (and Subramanian), whitelists and blacklists were well-known and used interchangeably long before the claimed priority date of the '040 Patent. Accordingly, instead of a blacklist, a POSA would have found it obvious to use a whitelist of IP addresses under the access authority of the host device in the Treuhaft combinations. As explained above and in more detail below in the Detailed Application section for the specific claim elements pertinent to the combination, a POSA would have had been motivated to combine Treuhaft with Sorenson and/or combine Treuhaft/Sorenson with Bellinson, and had a reasonable expectation in doing so.

85. Accordingly, as described above and below in the Detailed Discussion section, it is my opinion that the proposed combinations of Treuhaft, Treuhaft and Sorenson, and Treuhaft/Sorenson in view of Bellinson teach, or at least render obvious, "*discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list, wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device*," as recited by independent claims 1, 6, and 11. Thus, in my opinion, the proposed combinations of Treuhaft, Treuhaft and Sorenson, and Treuhaft/Sorenson in view of Bellinson present a substantial new question of patentability with respect to the Challenged Claims.

X. DETAILED APPLICATION OF THE PRIOR ART TO EVERY CLAIM FOR WHICH REEXAMINATION IS REQUESTED

86. I explain below in detail how, in my opinion, the prior art references listed above establish the unpatentability of the Challenged Claims as anticipated and obvious.

A. Ground 1: RFC 1928 in view of Koblas and RFC 3089 ("SOCKS") Render Obvious Claims 1, 4-6, and 9-11 of the '040 Patent

87. In my opinion RFC 1928 in view of Koblas and RFC 3089 ("SOCKS") render obvious claims 1, 4-6, and 9-11 of the '040 Patent.

1. Independent Claim 1

a. [1.pre] *"A packet receiving method, comprising:"*

88. To the extent the preamble is limiting, RFC 1928 describes SOCKS Protocol Version 5, which is a packet receiving method. The SOCKS Protocol is an Internet protocol that exchanges (i.e., sends and receives) network packets over TCP/IP between a client and server through a proxy server that acts as a firewall, called a SOCKS proxy server. SOCKS Version 5 provides authentication so only authorized users may access a server. According to RFC 1928, SOCKS Version 5 is a "protocol … designed to provide a framework for client-server applications in both the TCP and UDP domains to conveniently and securely use the services of a network firewall." Ex. 1005, 2.

89. Accordingly, in my opinion, RFC 1928 discloses "*a packet receiving method*," as claimed.

b. [1.1] *"receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request packet sent by the terminal device;"*

90. In my opinion, RFC 1928 and Koblas renders obvious this element. As laid out below, RFC 1928 discloses that a SOCKS proxy server receives a SOCKS connection request (claimed service request packet) sent by a SOCKS client (claimed terminal device). The SOCKS request carries a fully qualified domain name (claimed server domain name) indicating a server behind the SOCKS proxy server to which the SOCKS client desires a connection (claimed service server required by the service request packet). The SOCKS connection request also contains a source IP address of the SOCKS client (claimed terminal device)--rather than a terminal domain name. But, in my opinion, it would have been obvious to include a domain name in the SOCKS connection request based on the disclosure of Koblas, as I discuss below.

91. Specifically, RFC 1928 discloses that a SOCKS proxy server receives a SOCKS connection request (claimed service request packet) from a TCP-based SOCKS client (claimed terminal device):

> When a TCP-based client wishes to establish a connection to an object that is reachable only via a firewall (such determination is left up to the implementation), it must open a TCP connection to the appropriate SOCKS port on the SOCKS server system. The SOCKS service is conventionally located on TCP port 1080. If the connection request succeeds, the client enters a negotiation for the authentication method to be used, authenticates with the chosen method, then sends a relay request. The SOCKS server evaluates the request, and either establishes the appropriate connection or denies it.

Ex. 1005, 2-3; *see also id.*, 4 ("Once the method-dependent subnegotiation has completed, the client sends the request details.").

92. RFC 1928 shows the form of the SOCKS connection request:

"service request packet"
—SOCKS connection request—

The SOCKS request is formed as follows:

VER	CMD	RSV	ATYP	DST.ADDR	DST.PORT
1	1	X'00'	1	Variable	2

Where:

- VER protocol version: X'05'
- CMD
  - CONNECT X'01'
  - BIND X'02'
  - UDP ASSOCIATE X'03'
- RSV RESERVED
- ATYP address type of following address
  - IP V4 address: X'01'
  - DOMAINNAME: X'03'
  - IP V6 address: X'04'
- DST.ADDR desired destination address
- DST.PORT desired destination port in network octet order

Id., 4.* As highlighted above, the SOCKS connection request is a "packet" because it is a unit of information transmitted as a whole. Ex. 1012, 5 (definition of "packet"). Specifically, the SOCKS connection request is a TCP/IP packet sent over the Internet containing the VER, CMD, RSV, ATYP, DST.ADDR, and DST.PORT fields in a single message, as shown in the Figure above. Ex. 1005, 4; *see also id.*, 3. The SOCKS client is a TCP/IP client and thus communicates using TCP/IP packets. Accordingly, in my opinion, the SOCKS proxy server's receiving a SOCKS connection request from a SOCKS client discloses "receiving a service request packet," as claimed.

93. The SOCKS connection request "carries … a server domain name indicating a service server required by the service request packet sent by the terminal device," as claimed. Particularly, the SOCKS connection request carries a DST.ADDR (destination address) field containing a fully qualified domain name (claimed server domain name) of a server, behind the SOCKS proxy server, to which the SOCKS client desires to connect:

"service request packet"
—SOCKS connection request—

The SOCKS request is formed as follows:

VER	CMD	RSV	ATYP	DST.ADDR	DST.PORT
1	1	X'00'	1	Variable	2

"server domain name"
—fully-qualified domain name—

Where:

- VER protocol version: X'05'
- CMD
  - CONNECT X'01'
  - BIND X'02'
  - UDP ASSOCIATE X'03'
- RSV RESERVED
- ATYP address type of following address
  - IP V4 address: X'01'
  - DOMAINNAME: X'03'
  - IP V6 address: X'04'
- DST.ADDR desired destination address
- DST.PORT desired destination port in network octet order

Ex. 1005, 4.

94. As highlighted in the Figure above, the SOCKS request has a DST.ADDR field, which contains the "desired destination address" sought by the SOCKS client. The preceding field ATYP specifies the "address type of [the] following address" contained in the DST.ADDR field. *Id.*, 4. As shown, SOCKS supports a destination address in the form of "DOMAINNAME" using protocol version "X' 03'". *Id.*, 4. Following the description of the SOCKS connection request, RFC 1928 explains that X' 03' means "the address field contains a fully-qualified domain name." *Id.*, 5. Thus, in use, SOCKS clients send connection requests containing a fully qualified domain name (e.g., www.website.com) in the DST.ADDR field. Because the SOCKS connection request contains a fully qualified domain name of the destination server, it is my opinion that the SOCKS connection request discloses a "service request packet … carr[ying] … a server domain name indicating a service server required by the service request packet sent by the terminal device," as claimed.

95. The SOCKS connection request (claimed service request packet) carries a source IP address indicating the SOCKS client (claimed terminal device). TCP/IP protocol requires all TCP/IP packets to include a header containing a source IP address indicating the device that sent the packet. And because the SOCKS connection request is a TCP/IP packet sent over the Internet, it is required that the SOCKS connection request contains a header with the source IP address of the SOCKS client, though not shown in the Figure above. *Id.*, 1 (SOCKS is an Internet protocol), 2 (SOCKS is a framework for TCP-based client-server application), 7 (the connection established using SOCKS is a TCP connection). Thus, in my opinion, RFC 1928 discloses, or at least suggests, that the SOCKS connection request contains a source IP address indicating the SOCKS client (claimed terminal device).

i. Koblas discloses or suggests "*the service request packet carries a terminal domain name indicating the terminal device*"

96. To the extent that RFC 1928 does not expressly disclose that the SOCKS connection request carries a domain name indicating the SOCKS client (claimed terminal domain name), in my opinion Koblas suggests this by disclosing that the SOCKS server uses a Configuration File mapping SOCKS client domain names to corresponding permitted/denied server IP addresses. *See* Ex. 1006, 7-9.

> The configuration file is located on the firewall host and is used by sockd when determining whether to accept or deny requests. The file is parsed from beginning to end, with the first fully matching line returning the accessibility. The syntax of the lines in this file is as follows:
>
> {permit | deny} <source-host> <mask> [<<dest-host> <mask> [<operator> <port>]]
>
> source address destination address
>
> Lines begin with either 'permit' or 'deny' following which are either 2, 4, or 6 fields, containing host address and mask pairs for source and destination, as well as a boolean operator and a service port.

Ex. 1006, 7*.

97. As highlighted above, in the SOCKS Configuration File, the <source-host> field contains the host address of the SOCKS client source (claimed terminal device) and the corresponding <dest-host> field in the entry contains the address of the corresponding destination server. *Id.* Depending on which value the {permit | deny} field contains, the SOCKS server will respectively permit or deny the SOCKS client source in the <source-host> field to connect to the corresponding destination server address in the <dest-host> field. *Id.*, 7. Moreover, "[h]ost addresses and services may be specified either by name or number," meaning SOCKS supports listing either a domain name or an IP address in the <source-host> and <dest-host> fields. *Id.*,

98. Similarly, in the Sample Configuration file in Figure 5, Koblas shows the SOCKS server permits a connection request from the source domain name lloyd.mips.com (claimed terminal domain name) to connect to the corresponding preset destination address sgi.com.

FIGURE 5. A Sample Configuration File

```
#
# Deny all host to every host whois service
#
deny 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq whois
#
# Let lloyd.mips.com only use finger service to sgi.com
#            source                destination
permit lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0 eq finger
deny lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0
#
# Allow all hosts on the 130.62 network access to the world
#
permit 130.62.0.0 0.0.255.255
#
# Deny all hosts which do not match anything in this file
# (i.e. All hosts coming in from the Internet)
#
```

Id., 8*. In my opinion, because Koblas discloses that the SOCKS server screens connection requests based on the domain name of the SOCKS client source sending the connection request, Koblas at least suggests that the connection request may contain the domain name of the SOCKS client source. Thus, in my opinion, Koblas at least suggests "*the service request packet carries a terminal domain name indicating the terminal device*," as claimed. *Id.*

ii. Rational to combine Koblas with RFC 1928/RFC 3038

99. In my opinion, it would have been obvious based on Koblas to modify the SOCKS request to include the domain name of the client for several reasons. For example, Koblas provides teaching, suggestion, and/or motivation for making this modification. Whereas RFC 1928 describes SOCKS generally to "provide a framework for client-server applications … to conveniently and securely use the services of a network firewall," Ex. 1005, 2, Koblas drills down with "several strategies which can be used to configure an Internet connection to prevent unwanted intrusion" using the SOCKS protocol, Ex. 1006, 1. Thus, in my opinion, a POSA considering RFC 1928 and the general framework of the SOCKS protocol would have looked to Koblas for specific examples of strategies to configure secure connections using SOCKS.

100. As I explain above, Koblas teaches that the SOCKS protocol applies access controls to SOCKS connection requests using information contained in a "Configuration File." *Id.*, 7-9. The Configuration File maps terminal domain names to corresponding allowed or denied server addresses. *See Id.*, 8 (for example, permitting the terminal domain name "lloyd.mips.com" to access the service server domain name "sgi.com"). A POSA would have understood this means that, in at least some scenarios, the SOCKS server needs the domain name—not the IP address—of the SOCKS client to determine whether to permit or deny the connection request. There are two ways the SOCKS server can determine the domain name of the SOCKS client:

101. The first, and most logical, way is that the SOCKS client simply provides its domain name to the SOCKS server in the connection request. The SOCKS server needs the SOCKS client's domain name to apply the Configuration File rules for that SOCKS client to its connection request, and a POSA would have recognized having the SOCKS client simply include its domain name in a connection request as the simplest and most apparent way for the SOCKS server to learn domain name of the SOCKS client. For this reason, it is my opinion that a POSA would have found it obvious modify SOCKS as described in RFC 19298 to include the domain name of the SOCKS client in the connection request.

102. The second way, the SOCKS server resolves the domain name of the SOCKS client from its source IP address (e.g., contained in the connection request packet). And, after resolving the domain name of the SOCKS client, the SOCKS server may then identify and apply that client's rules in the Configuration File. But this way is less efficient than the first way, requiring the extra step of resolving the SOCKS client domain name. Accordingly, Koblas suggests that, in at least

some implementations, the SOCKS proxy server would need to convert the source IP address of the SOCKS client, contained in the SOCKS connection request, to its corresponding domain name before applying the access controls specified in the Configuration File.

103. To eliminate this extra step of converting the IP address of the SOCKS client to a domain name before applying the Configuration File, in my opinion, it would have been obvious to simply have the SOCKS client include its domain name in the SOCKS connection request (rather than just its IP address). This would allow the SOCKS proxy server to immediately apply the access controls in the Configuration File upon receiving a connection request, speeding up and simplifying the process of establishing the connection. It would also eliminate the need for the SOCKS proxy server to maintain and update information mapping SOCKS client IP addresses to domain names, request services from other devices to convert client IP addresses to domain names, and/or perform other DNS-like functions to convert the source IP address into a domain name so the Configuration File can be applied. Thus, the combination would allow simplifying the configuration of the DNS proxy server itself could and streamlining the process for establishing a secure connection. Accordingly, in my opinion, a POSA would have been motivated to modify the SOCKS protocol as described in RFC 1928 to include the domain name of the SOCKS client in a connection request.

104. In my opinion, a POSA would have made this modification with mere routine skill in the art and a reasonable expectation of success. Indeed, a fully qualified domain name is essentially equivalent to its corresponding IP address because both represent the same address of the same device on a network. Ex. 1012 (defining "domain name" as "[a]n address of a network connection that identifies the owner of that address in a hierarchical format"). They differ mainly in their form: a domain name is easily read and memorized by a person whereas a machine-readable IP address is not. The '040 Patent itself recognizes this essential equivalence between a domain name and its IP address, explaining that "mutual conversion [can be achieved] between a domain name which is readily memorized by a user and a machine recognizable IP address." Ex. 1001, 8:30-33.

105. Accordingly, in my opinion, a POSA would have viewed modifying the SOCKS connection request to include the client's domain name---when the connection request already contains that same address in its other form as an IP address—as a simple, unintrusive change. For

example, a POSA would have expanded the structure of the SOCKS connection request, discussed above, to include an additional client domain name field containing the domain name of the SOCKS client. This minor change would not otherwise impact operation of SOCKS protocol and, as discussed, advantageously simplifies the process of applying the access controls in the Configuration File and related functionality.

106. Accordingly, in my opinion, RFC 1928 and Koblas render obvious "receiving a service request packet sent by a terminal device, wherein the service request packet carries … a server domain name indicating the service server requested by the service request packet sent by the terminal device," as claimed.

c. [1.2] *"resolving the received server domain name to obtain a service server Internet protocol (IP) address; and"*

107. In my opinion, RFC 1928 and RFC 3089 teach or at least suggest this element. Resolving the received fully qualified domain name of the server (claimed received server domain name) to obtain an IP address of that server (claimed service server IP address) is a basic and necessary function of SOCKS. Accordingly, although RFC 1928 does not explicitly detail this process, it is my opinion that a POSA would understand it to be taught, or at least suggested, by RFC 1928's disclosure and teaching of the SOCKS protocol. *See* RFC 3089. RFC 1928 only lacks a detailed description on this aspect because RFC 1928 focuses on the broader framework of the SOCKS protocol and does not attempt to describe such well-known and basic aspects of the protocol. *See generally* Ex. 1005. Nevertheless, RFC 3089 does explicitly teach resolving the IP address of the server domain name as claimed in this limitation. Thus, in my opinion, it would have been obvious to do so in view of the teachings of RFC 1928 and RFC 3089.

i. RFC 3089 discloses element [1.2]

108. RFC 3089 describes a gateway mechanism for the SOCKS proxy server to handle both IPv6 and IPv4. *See* Ex. 1007, 1 ("The SOCKS-based IPv6/IPv4 gateway mechanism is based on a mechanism that relays two 'terminated' IPv4 and IPv6 connections at the 'application layer' (the SOCKS server)"). RFC 3098 explains that the "characteristics [of the SOCKS server] are inherited from those of the connection relay mechanism at the application layer and those of the native SOCKS mechanism." *Id.* 1. Thus, a POSA would have understood that the SOCKS server discussed in RFC 3089 inherits, and thus includes, the functionality of the SOCKS server described

in the earlier RFC 1928. That is, when RFC 3089 refers to the SOCKS server, it is the same SOCKS server referenced in RFC 1928.

109. RFC 3089 describes in detail the process by which the SOCKS server resolves the fully qualified domain name (FQDN) of the destination node (Destination D)—the claimed service server—to obtain its "real IP address"—the claimed service server IP address:

> The detailed internal procedure of the "DNS name resolving delegation" and address mapping management related issues are described as follows.
>
> 1. An application on the source node (Client C) tries to get the IP address information of the destination node (Destination D) by calling the DNS name resolving function (e.g., gethostbyname()). At this time, the logical host name ("FQDN") information of the Destination D is passed to the application's *Socks Lib* as an argument of called APIs.
> 2. Since the *Socks Lib* has replaced such DNS name resolving APIs, the real DNS name resolving APIs is not called here. The argued "FQDN" information is merely registered into a mapping table in *Socks Lib*, and a "fake IP" address is selected as information that is replied to the application from a reserved special IP address space that is never used in real communications (e.g., 0.0.0.x). The address family type of the "fake IP" address must be suitable for requests called by the applications. Namely, it must belong to the same address family of the Client C, even if the address family of the Destination D is different from it. After the selected "fake IP" address is registered into the mapping table as a pair with the "FQDN", it is replied to the application.
> 3. The application receives the "fake IP" address, and prepares a "socket". The "fake IP" address information is used as an element of the "socket". The application calls socket APIs (e.g., connect()) to start a communication. The "socket" is used as an argument of the APIs.
> 4. Since the *Socks Lib* has replaced such socket APIs, the real socket function is not called. The IP address information of the argued socket is checked. If the address belongs to the special address space for the fake address, the matched registered "FQDN" information of the "fake IP" address is obtained from the mapping table.
> 5. The "FQDN" information is transferred to the *Gateway* on the relay server (Gateway G) by using the SOCKS command that is matched to the called socket APIs. (e.g., for connect(), the CONNECT command is used.)
> 6. Finally, the real DNS name resolving API (e.g., getaddrinfo()) is called at the *Gateway*. At this time, the received "FQDN" information via the SOCKS protocol is used as an argument of the called APIs.
> 7. The *Gateway* obtains the "real IP" address from a DNS server, and creates a "socket". The "real IP" address information is used as an element of the "socket".
> 8. The *Gateway* calls socket APIs (e.g., connect()) to communicate with the Destination D. The "socket" is used as an argument of the APIs.

Ex. 1007, 5-6. Accordingly, in my opinion, RFC 3089 discloses "*resolving the received server domain name to obtain a service server Internet protocol (IP) address*," as claimed.

ii. Rationale to Combine RFC 3089 with RFC 1928

110. In my opinion, a POSA would have found it obvious to modify RFC 1928 to resolve the received fully qualified domain name of the server (claimed received server domain name) to obtain an IP address of that server (claimed service server IP address), as disclosed by RFC 1928. First, the references provide teaching, suggestion, and/or motivation for making this modification. Both RFC 1928 and RFC 3089 relate to the same SOCKS protocol, so a POSA considering the general framework of SOCKS in RFC 1928 would have looked to other SOCKS references like RFC 3089 for additional implementation details about various aspects of the system.

111. Moreover, at the outset of describing SOCKS' domain name resolution process, RFC 3089 explains that, "[i]n all communication applications, it is [] necessary to obtain destination IP address information to start a communication." Ex. 1007, 4. As explained above, RFC 1928 teaches that SOCKS supports receiving the destination address as a fully qualified domain name in the connection request. Ex. 1005, 4-5. Thus, according to RFC 1928, it is absolutely necessary in this scenario to resolve the fully qualified domain name into a corresponding IP address—or else the SOCKS server cannot start the connection. Ex. 1007, 4. Accordingly, in my opinion, to implement a working SOCKS system, a POSA would have found it obvious (and necessary) to modify SOCKS as described in RFC 1928 to resolve the fully qualified domain name of the destination server (claimed received server domain name) to obtain an IP address of the destination server (claimed service server IP address), as claimed.

112. Moreover, the combination merely involves the use of known technique (RFC 3089's resolving a fully qualified domain name) to improve the same system (SOCKS communication system according to RFC 1928) in the same way, or applying a known technique (RFC 3089's resolving a fully qualified domain name) to a known system (SOCKS communication system according to RFC 1928) ready for improvement to yield predictable results. For example, since the SOCKS protocol as described in RFC 1928 would be unable to establish a connection without first resolving the fully qualified domain name of the destination as described in RFC 3089, adding this functionality would improve SOCKS as described in RFC 1928 in the same way it operates in RFC 3089. Additionally, because the combination simply adds necessary SOCKS functionality to the general SOCKS system described in RFC 1928 to make it operational, the combination would yield predictable results.

113. Accordingly, in my opinion, RFC 1928 and RFC 3089 render obvious "resolving the received server domain name to obtain a service server Internet protocol (IP) address," as claimed.

d. [1.3] *"discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list,"*

114. In my opinion, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious this element. Koblas discloses or suggests discarding the SOCKS connection request (service request packet) if the resolved address of the destination server (resolved service server IP address) is not listed as an allowed address (does not belong to a preset service server IP address) corresponding to the domain name of the SOCKS client (received terminal domain name) in a Configuration File (preset list). Moreover, in my opinion, it would have been obvious to modify the RFC 1928/RFC 3089 combination discussed above to include this element.

i. Koblas discloses element [1.3]

115. RFC 1928 discloses that the SOCKS server discards the SOCKS connection request (claimed service request packet) if the requested connection is not "appropriate": [t]he SOCKS server evaluates the request, and either establishes the appropriate connection or denies it." Ex. 1005, 3. As explained above, the broadest reasonable interpretation of "discarding the service request packet" includes preventing unauthorized access to the resolved service server IP address. *Supra* Section VI.A In my opinion, denying the connection request, as disclosed by RFC 1928, constitutes preventing unauthorized access to the resolved service server IP address.

116. RFC 1928, however, does not expressly disclose the mechanism the SOCKS server uses to determine whether the requested connection is "appropriate." Thus, RFC 1928 does not expressly disclose discarding the SOCKS connection request "if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list," as claimed.

117. Koblas, however, teaches that the SOCKS server uses a "Configuration File" (claimed preset list) to evaluate and allow or deny the connection request. The Configuration File contains an entry for each SOCKS client source (claimed terminal device) identifying

corresponding "permit" or "deny" destination addresses (claimed service server IP address) to which the SOCKS server will respectively permit or deny connections requested by the SOCKS client source:

> The configuration file is located on the firewall host and is used by sockd when determining whether to accept or deny requests. The file is parsed from beginning to end, with the first fully matching line returning the accessibility. The syntax of the lines in this file is as follows:
>
> 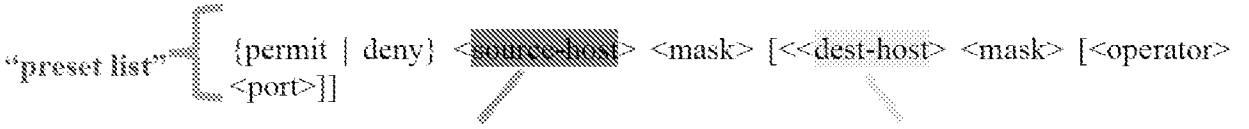
>
>
> "terminal domain name" "corresponding preset service server IP address"
>
> Lines begin with either 'permit' or 'deny' following which are either 2, 4, or 6 fields, containing host address and mask pairs for source and destination, as well as a boolean operator and a service port.

Ex. 1006, 7*.

118. As highlighted above, in the SOCKS Configuration File, the <source-host> field contains the host address of the SOCKS client source (claimed terminal device) and the corresponding <dest-host> field in the entry contains the address of the corresponding destination server (claimed corresponding preset service server IP address). *Id.* Depending on which value the {permit | deny} field contains, the SOCKS server will respectively permit or deny the SOCKS client source in the <source-host> field to connect to the corresponding destination server address in the <dest-host> field. *Id.*, 7. Additionally, Koblas explains that "[h]ost addresses and services may be specified either by name or number," meaning SOCKS supports listing either a domain name or an IP address in the <source-host> and <dest-host> fields. *Id.*, 8. Accordingly, in my opinion, the Configuration File of Koblas discloses the claimed preset list containing a preset service server IP address <dest-host> corresponding to the received terminal domain name <source-host>.

119. Koblas further teaches that the SOCKS server discards the SOCKS connection request if the Configuration File (claimed preset list) does not list the <dest-host> address (claimed resolved service server IP address) as a permitted connection for the <source-host> domain name (claimed does not belong to a preset service server IP address corresponding to the received terminal domain name). This is because "[a]ccess is denied to all addresses which do not match

anything in the configuration file." *Id.*, 8. Thus, if the SOCKS server does not find a particular resolved <dest-host> address listed in the Configuration File as permitted connection for the <source-host> domain name, it will deny the SOCKS connection request even though the Configuration File does not list the <dest-host> address as a denied connection. *Id.*, 8.

120. In Figure 5, Koblas "shows an example of how the lines in a configuration file might appear":

"preset list"

FIGURE 5. A Sample Configuration File

```
#
# Deny all host to every host whois service
#
deny 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq whois
#
# Let lloyd.mips.com only use finger service to sgi.com
#    "terminal domain name"    "preset service server IP address"
permit lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0 eq finger
deny lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0
#
# Allow all hosts on the 130.62 network access to the world
#
permit 130.62.0.0 0.0.255.255
# "discarding ... if the resolved service server IP address does not belong ..."
#
# Deny all hosts which do not match anything in this file
# (i.e. All hosts coming in from the Internet)
#
```

Id., 8. In this Sample Configuration file, the SOCKS server permits a request from the source domain name lloyd.mips.com (claimed terminal domain name) to connect to the corresponding preset destination address sgi.com, and denies requested connections "which do not match anything in this file." *Id.* Although shown as a domain name in this example, the server destination address sgi.com "may be specified either by name or number," so the IP address of sgi.com could be used instead. *Id.*, 8;

121. As explained above, "discarding" the service request includes preventing unauthorized access to the resolved service server IP address. *Supra* Section VI.A. In my opinion, by allowing a connection if the Configuration File lists the resolved server IP address as an allowed connection for the source domain name and denying all connections that do not match anything in

the Configuration File, application of the Configuration File prevents unauthorized access to the resolved server IP address.

122. Accordingly, for at least the reasons above, it is my opinion that Koblas discloses "discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list," as claimed.

ii. Rationale to combine Koblas with RFC 1928 and RFC 3089

123. In my opinion, a POSA would have found it obvious to modify SOCKS as described in RFC 1928 and/or RFC 3089 to use a Configuration File as described in Koblas to discard a SOCKS connection request if the Configuration File does not list the resolved destination address as an allowed connection for the SOCKS client. The references provide teaching, suggestion, and/or motivation for making this modification. RFC 1928, RFC 3089, and Koblas all relate to the same SOCKS protocol, so a POSA implementing a SOCKS system as described in RFC 1928/RFC 3089 would have looked to other SOCKS references like Koblas for additional implementation details about various aspects of the system.

124. Although RFC 1928 explains "[t]he SOCKS server evaluates the request, and either establishes the appropriate connection or denies it," RFC 1928 does not discuss a specific mechanism within SOCKS to do so because this falls outside its general scope of "provid[ing] a framework for client-server applications … to conveniently and securely use the services of a network firewall." Ex. 1005, 2. Thus, in my opinion, a POSA would have looked to other SOCKS references for specific mechanisms to evaluate the appropriateness of SOCKS connection requests, and Koblas's Configuration File provides one such mechanism.

125. Koblas addresses "[o]ne of the more important [security] issues" to consider when connecting to a network over the Internet: "intruders attempting to gain access to local hosts" using the SOCKS protocol. Ex. 1006, 3. In my opinion, a POSA would have found this solution pertinent to RFC 1928 because, in applying the Configuration File as taught by Koblas, the SOCKS server performs the function mentioned in RFC 1928: "[t]he SOCKS server evaluates the request, and either establishes the appropriate connection or denies it." Ex. 1005, 3. Thus, a POSA would have found it obvious and been motivated to add the SOCKS Configuration File functionality described

in Koblas to the SOCKS system as described in RFC 1928 to make SOCKS work as RFC 1928 intends.

126. Moreover, the combination merely involves the use of a known technique (Koblas's Configuration File) to improve the same system (SOCKS system according to RFC 1928) in the same way, or applying a known technique (Koblas's Configuration File) to a known system (SOCKS system according to RFC 1928) ready for improvement to yield predictable results. The combination would improve SOCKS as described in RFC 1928 because, as discussed, RFC 1928 teaches that the SOCKS server evaluates a connection request to confirm its appropriateness before establishing the connection, but does not expressly describe a mechanism for doing so. Ex. 1005, 2. Koblas fills this gap with its Configuration-File solution for evaluating SOCKS connection requests. Thus, combining Koblas's Configuration File technique with RFC 1928 would improve the SOCKS system as described in RFC 1928 in the same way it works in Koblas—preventing intruders from gaining access to local hosts by only allowing connections from certain sources to certain destinations.

127. Additionally, in my opinion, the combination would yield predictable results and would be made through only routine skill in the art. Indeed, the combination simply adds a SOCKS security solution from Koblas to the general SOCKS system described in RFC 1928. With the Configuration File functionality added to RFC 1928, the combined SOCKS system would operate in the way RFC 1928 describes: the SOCKS server would evaluate a connection request and establish the connection if appropriate or deny it otherwise. Ex. 1005, 3. Moreover, Koblas explains, "[t]he configuration file is located on the firewall host and is used by sockd when determining whether to accept or deny requests." Ex. 1006, 7. By firewall, Koblas is referring to the SOCKS server. Thus, a POSA would have recognized that, in combining Koblas with RFC 1928, the Configuration File of Koblas would be stored on the SOCKS server in RFC 1928 and applied when a connection request is received. Because Koblas and RFC 1928 both describe the same SOCKS protocol, in my opinion, a POSA would have modified the SOCKS server in RFC 1928 to include the Configuration-File functionality of Koblas through routine skill had and a reasonable expectation of success in doing so. In fact, in my experience, most, if not all, SOCKS implementations available at the time and to this date use a configuration file that provides similar if not identical features as disclosed by Koblas.

128. Accordingly, for at least the reasons above, it is my opinion that RFC 1928, Koblas, and RFC 3089 render obvious "discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list," as claimed.

e. [1.4] *"wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device."*

129. In my opinion, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious this element. As highlighted below, in Koblas's Configuration File (preset list), the domain name of each SOCKS client address (terminal domain name) is provided with an accessible destination server address (accessible service server IP address) under an access authority of the SOCKS client source (terminal device):

{permit | deny} <source-host> <mask> [<dest-host> <mask> [<operator> <port>]]

"terminal domain name" "accessible service server IP address"

Ex. 1006, 7.

"preset list"

FIGURE 5. A Sample Configuration File

```
#
# Deny all host to every host whois service
#
deny 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq whois
#
# Let lloyd.mips.com only use finger service to sgi.com
#    "terminal domain name"    "accessible service server IP address"
permit lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0 eq finger
deny lloyd.mips.com 0.0.0.0 sgi.com 0.0.0.0
#
# Allow all hosts on the 130.62 network access to the world
#                    "accessible service server IP address"
permit 130.62.0.0 0.0.255.255
#"terminal domain name"
# Deny all hosts which do not match anything in this file
# (i.e. All hosts coming in from the Internet)
#
```

Ex. 1006, 8. As explained above, the Configuration File can specify source and destination addresses "either by name or number." *Id.*, 6. Thus, although the Sample Configuration File above shows an IP address (e.g., 130.62.0.0) for some of the SOCKS client sources (claimed terminal device), Koblas teaches that a domain name could be used instead. Thus, in my opinion, the combination renders obvious "the terminal domain name of each terminal device" is provided with a corresponding accessible service server IP address, as claimed.

130. Moreover, though Koblas only shows one accessible (i.e., permitted) destination address for each SOCKS client, it would have been obvious to include additional "permit" entries listing additional accessible IP addresses for each SOCKS client. For example, the Sample Configuration File only permits the SOCKS client lloyd.mips.com to access a single server address—sgi.com. Of course, in practice, this SOCKS client and/or its user may need to access more than just one site on the network. In this case, it would have been obvious for the administrator to add more "permit" entries to the Configuration file listing more accessible server addresses for the SOCKS client lloyd.mips.com.

131. Accordingly, for at least the reasons above, it is my opinion that RFC 1928, Koblas, and RFC 3089 render obvious "wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device," as claimed.

2. Dependent Claim 4

a. [4.1] ***"The method according to claim 1, wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:***

if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device."

132. I understand that claim 4 depends on independent claim 1 and recites the additional elements quoted above. In my opinion, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious independent claim 1 for the reasons discussed above. *Supra* Section X.A.1.

Moreover, it is my opinion that the combination renders obvious the additional elements of claim 4.

133. For example, RFC 3038 discloses "[t]he SOCKS server evaluates the request, and either establishes the appropriate connection or denies it." Ex. 1005, 3. Additionally, as discussed, in the combined SOCKS system, the SOCKS server establishes a connection request if the Configuration File lists the resolved IP address of the destination server as a "permitted" connection for the domain name of the particular SOCKS client making the connection request. *See, e.g.*, Ex. 1006, 7-8. Accordingly, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious "*wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises: if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device*," as recited in claim 4.

3. Dependent Claim 5

a. [5.1] "*The method according to claim 1, wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:*

if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device."

134. Claim 5 depends on independent claim 1 and recites the additional elements quoted above. The combination of RFC 1928, Koblas, and RFC 3089 renders obvious independent claim 1 for the reasons discussed above. *Supra* Section X.A.1. Moreover, in my opinion, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious additional elements of claim 5.

135. For example, in Koblas, the Configuration File specifies that the SOCKS client domain name lloyd.mips.com may access the server sgi.com, but only for purposes of using the "finger service" provided by that server. Ex. 1006, 8. This is reflected in a comment in the

Configuration File above the permit and deny entries for lloyd.mips.com, stating "Let lloyd.mips.com only use finger service to sgi.com." *Id.* Accordingly, the "permit" entry for lloyd.mips.com specifies that lloyd.mips.com can only use sgi.com's service "eq finger", and the "deny" entry beneath it denies all other connections sought by lloyd.mips.com to the server sgi.com. *Id.*, 8.

136. Thus, in my opinion, when the SOCKS client lloyd.mips.com sends a connection request for sgi.com, the SOCKS server applies the Configuration File to determine that sgi.com is an accessible address for lloyd.mips.com. *Id.* In my opinion, this discloses determining "if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list," as claimed. If this is the case, the SOCKS server then determines that the SOCKS client is seeking access only to the "finger service" of sgi.com, and not to the server more broadly, before establishing the connection. *Id.* In my opinion, this discloses "determining a service type of the service request according to the terminal domain name of the terminal device," as claimed.

137. Accordingly, in my opinion, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious "wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises: if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device," as recited in claim 5.

4. Independent Claim 6

138. Independent claim 6 is an apparatus claim to a "deep packet inspection (DPI) device" configured to perform a method including elements [6.pre] to [6.4], which are virtually identical to respective elements [1.pre] to [1.4] of independent claim 1. *Compare* '040 Patent claim 1 *with id.*, claim 6. As set forth below, in my opinion, the combination of RFC 1928, Koblas, and RFC 3089 renders the elements of independent claim 6 for reason similar to those discussed above for independent claim 1.

a. [6.pre] ***"A deep packet inspection (DPI) device comprising a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method comprising:"***

139. To the extent the preamble is limiting, the SOCKS server described in RFC 1928, Koblas, and RFC 3089 discloses the claimed DPI device. Moreover, since a server is a computer, it would be understood to have a non-transitory computer readable storage medium (e.g., memory), instructions stored in memory for performing the functions attributed to the SOCKS server in the references, and one or more processors that execute the instructions to perform those functions.

140. Accordingly, in my opinion, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious "[a] deep packet inspection (DPI) device comprising a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method," as claimed.

b. [6.1] *"receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request packet sent by the terminal device;"*

141. *See* the discussion above for element [1.1]. *Supra* Section X.A.1.b.

c. [6.2] *"resolving the server domain name to obtain a service server Internet protocol (IP) address; and"*

142. *See* the discussion above for element [1.2]. *Supra* Section X.A.1.c.

d. [6.3] *"discarding the packet if the service server IP address resolved does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list,"*

143. *See* the discussion above for element [1.3]. *Supra* Section X.A.1.d.

e. [6.4] *"wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with accessible service server IP addresses under an access authority of the terminal device."*

144. *See* the discussion above for element [1.4]. *Supra* Section X.A.1.e.

5. Dependent Claim 9

a. [9.1] *"The DPI device according to claim 6, wherein after the*

resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:

if the service server IP address resolved belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device."

145. Claim 9 depends on independent claim 6 and recites the additional elements quoted above. The combination of RFC 1928, Koblas, and RFC 3089 renders obvious independent claim 6 for the reasons discussed above. *Supra* Section X.A.4. Moreover, claim 9 recites subject matter virtually identical to that discussed above for dependent claim 4. Accordingly, for the reasons discussed above in connection with claim 4, it is my opinion that the combination of RFC 1928, Koblas, and RFC 3089 renders obvious the subject matter of claim 9. *Supra* Section X.A.2

6. Dependent Claim 10

a. [10.1] ***"The DPI device according to claim 6, wherein after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:***

if the service server IP address resolved belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device."

146. Claim 10 depends on independent claim 6 and recites the additional elements quoted above. The combination of RFC 1928, Koblas, and RFC 3089 renders obvious independent claim 6 for the reasons discussed above. *Supra* Section X.A.4. Moreover, claim 10 recites subject matter virtually identical to that discussed above for dependent claim 5. Accordingly, for the reasons discussed above in connection with claim 5, it is my opinion that the combination of RFC 1928, Koblas, and RFC 3089 renders obvious the subject matter of claim 10. *Supra* Section X.A.3.

7. Independent Claim 11

147. Independent claim 11 is a system claim to a "deep packet inspection (DPI) device" and a "terminal device." As explained above, the SOCKS server and SOCKS client in RFC 1928,

Koblas, and RFC 3089 respectively correspond to the claimed DPI device and terminal device. *Supra* Section X.A.4. Substantively, claim 11 is virtually identical to independent claims 1 and 6 discussed above. Accordingly, as set forth below, it is my opinion that the combination of RFC 1928, Koblas, and RFC 3089 renders obvious the elements of independent claim 11 for reasons similar to those discussed above for independent claims 1 and 11.

a. [11.pre] ***"A system, comprising: a deep packet inspection (DPI) device; and a terminal device:"***

148. To the extent the preamble is limiting, the combination of RFC 1928, Koblas, and RFC 3089 renders obvious this element. *Supra* Sections X.A.4.a, 1.a.

b. [11.1] ***[a terminal device]*** ***"configured to send a service request packet to the DPI device, wherein the packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request sent by the terminal device;"*** **and**

"the DPI device having a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method comprising: receiving the service request packet sent by the terminal device;"

149. *See* the discussion above for element [1.1]. *Supra* Section X.A.1.b.

c. [11.2] ***"resolving the server domain name to obtain a service server Internet protocol (IP) address; and"***

150. *See* the discussion above for element [1.2]. *Supra* Section X.A.1.c.

d. [11.3] ***"discarding the packet if the service server IP address resolved does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list,"***

151. *See* the discussion above for element [1.3]. *Supra* Section X.A.1.d.

e. [11.4] ***"wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with accessible service server IP addresses under an access authority of the terminal device."***

152. *See* the discussion above for element [1.4]. *Supra* Section X.A.1.e.

B. Grounds 2 and 3: Treuhaft Anticipates and/or Renders Obvious Claims 1, 4-6, and 9-11 of the '040 Patent

153. In my opinion, Treuhaft anticipate and/or renders obvious claims 1, 4-6, and 9-11 of the '040 Patent.

1. Independent Claim 1

a. [1.pre] *"A packet receiving method, comprising:"*

154. To the extent the preamble is limiting, Treuhaft discloses a packet receiving method. Generally, in Treuhaft, a host device 105 sends a DNS query 110 to a DNS name server 120. Ex. 1008, ¶¶ [0025]-[0027], [0030]-[0032]. The DNS query 110 requests the DNS name server 102 to resolve a domain name, contained in the DNS query, into a corresponding IP address and return the resolved IP address to the host device 105. *Id.* In response to receiving the DNS query 110, the DNS name server 102 performs the steps shown in FIGS. 5B and 5C to resolve the domain name into a corresponding IP address and return the IP address to the host device 105 in a DNS response 170. *Id.*, ¶¶ [0027], [0032], [0064]-[0067]. The claimed terminal device, packet(s), and packet receiving method are highlighted in Figures 1, 5B, and 5C of Treuhaft below:

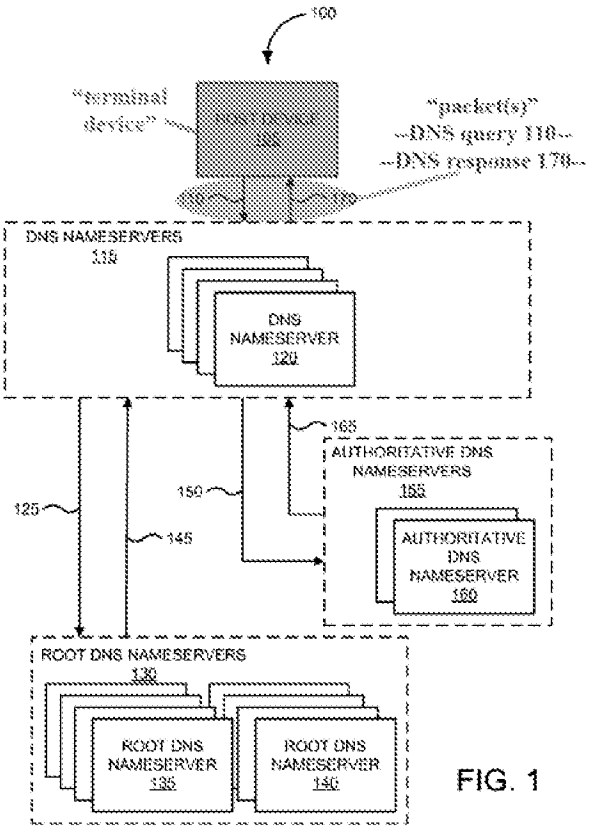

FIG. 1

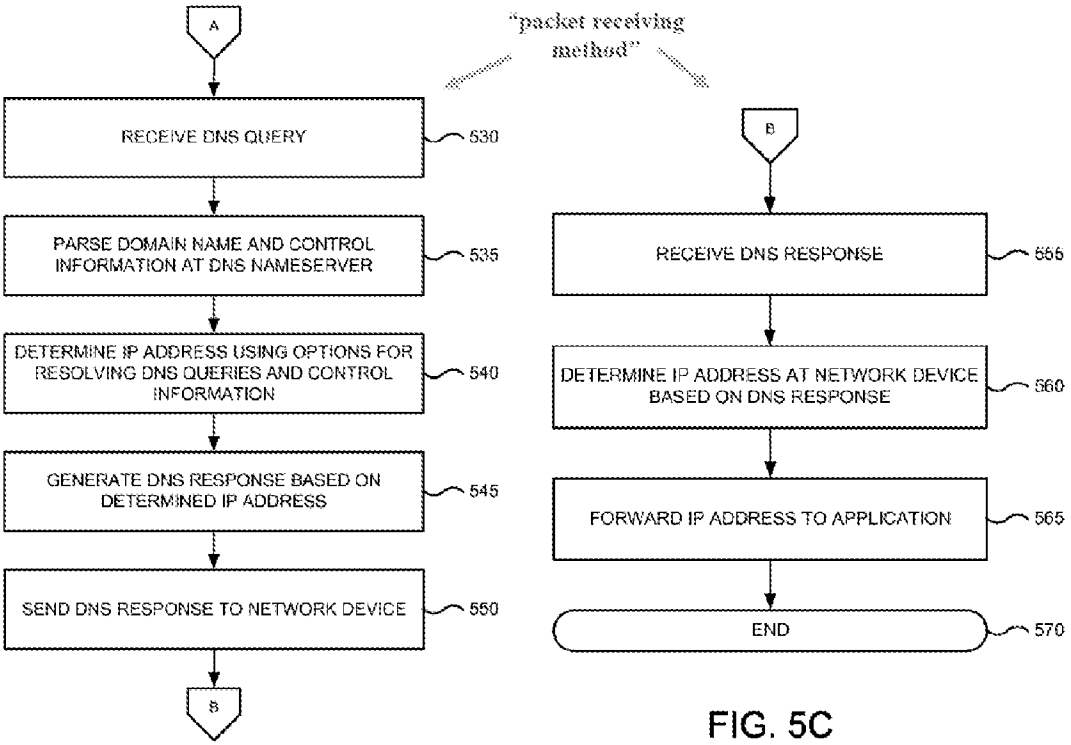

FIG. 5B

FIG. 5C

Treuhaft, FIGS. 1, 5B, 5C*

155. Accordingly, in my opinion, Treuhaft discloses "*a packet receiving method*," as claimed.

b. [1.1] *"receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request packet sent by the terminal device;"*

156. In my opinion, Treuhaft discloses, or at least suggests, this element. Specifically, as explained below, Treuhaft discloses receiving a service request packet (modified DNS query 110) sent by a terminal device (host device 105). Additionally, the service request packet (modified DNS query 110) carries a terminal domain name (control information) indicating the terminal device (host device 105) and a server domain name (domain name) indicating the service server requested by the service request packet (modified DNS query 110) sent by the terminal device (host device 105).

157. In step 525 of FIG. 5A, the host device 105 sends a modified DNS query to the DNS name server 120, and the DNS name server 120 receives the modified DNS query in step 530 of FIG. 5B. Ex. 1008, ¶ [0063] ("In step 525, the modified DNS query is sent to a DNS nameserver. For example, the modified DNS query may be sent to DNS nameserver 120."), ¶ [0064] ("in step 530, the DNS query is received"). In my opinion, Treuhaft's DNS name server 120 receiving the modified DNS query from the host device 105 corresponds to the claimed receiving a service request packet sent by a terminal device:

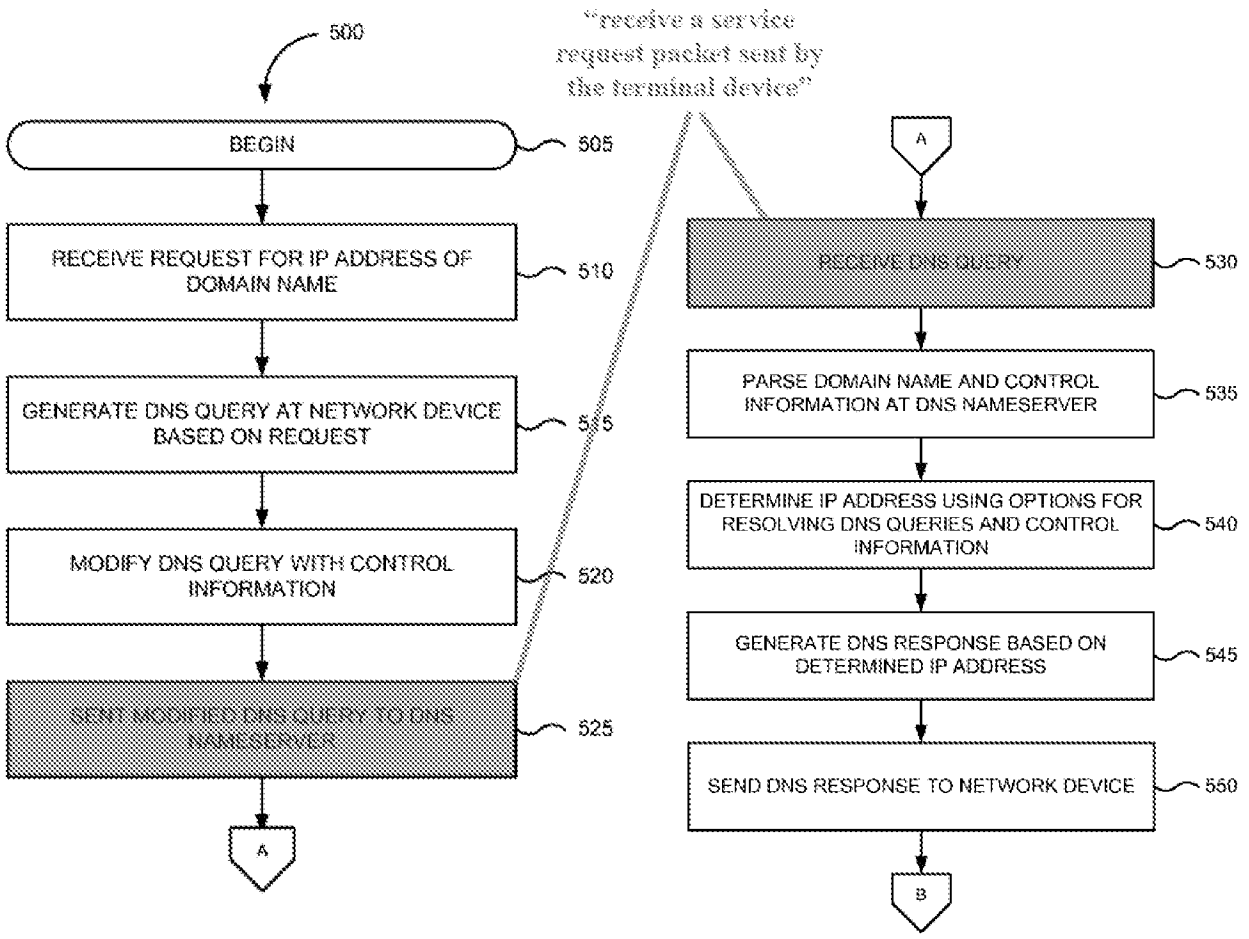

FIG. 5A FIG. 5B

158. Treuhaft's DNS modified DNS query discloses "a service request packet … carr[ying] … a server domain name indicating a service server required by the service request packet sent by the terminal device," as claimed. Treuhaft explains that "host device 105 makes DNS query 110, for example for the IP address of the domain name 'www.cnet.com,' to a set of DNS nameservers 115." *Id.*, ¶ [0025]. That is, in this example, the DNS query contains the domain name "www.cnet.com" of a server host device 105 seeks to access, and host device 105 needs the server's IP address from the DNS name server to do so. In Figure 2, Treuhaft shows an example format of the DNS query 400:

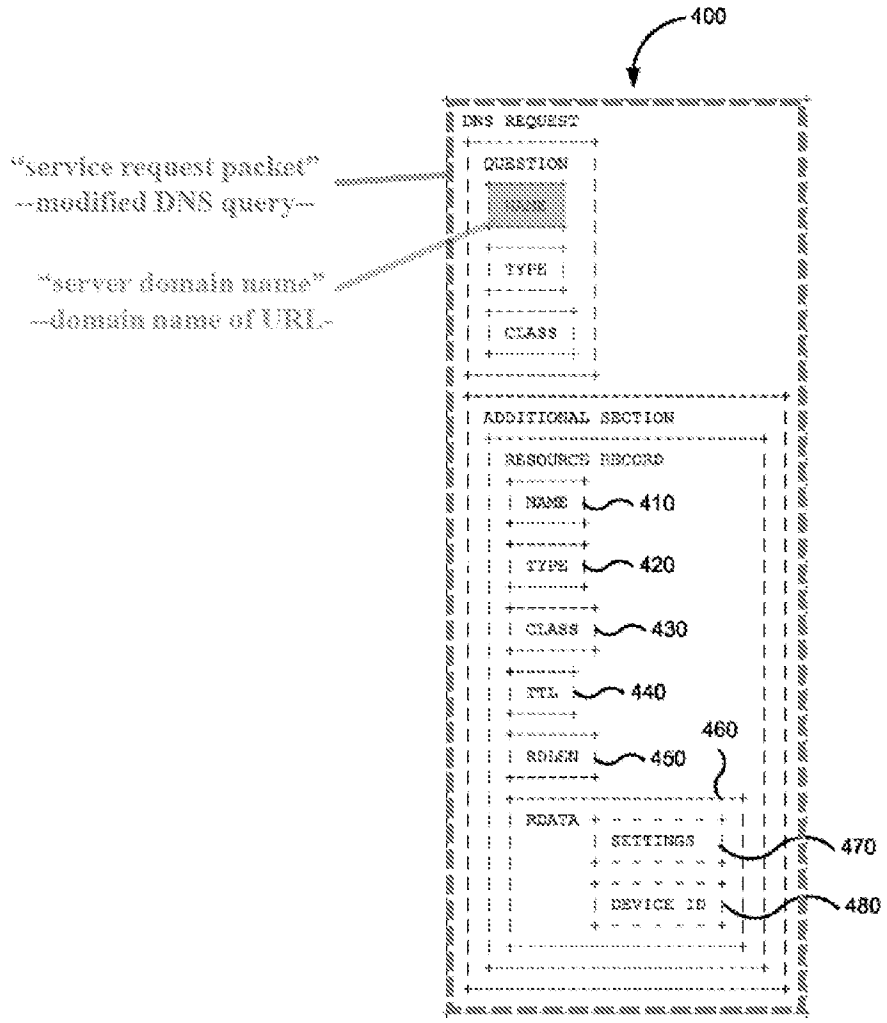

Treuhaft, FIG. 4*

As highlighted in Figure 4 above, the DNS query 400 includes a NAME field containing the domain name of the URL the host device 105 seeks to access. *Id.*, ¶¶ [0054]. Thus, like the '040 Patent's service request packet, the DNS query of Treuhaft "carries … a server domain name indicating a service server required by the service request." Ex. 1001, 3: 27-31; *see also id.*, FIG. 4 (illustrating the claimed service request packet as an HTTP request containing a domain name).

159. Under the broadest reasonable interpretation standard, the modified DNS query of Treuhaft is a "packet" because is a unit of information transmitted as a whole: a DNS protocol request message. *See id.*, ¶¶ [0041], [0043] (alternatively referring to the DNS query as a "DNS message"); Ex. 1012, 5 (definition of "packet"). A DNS query is a self-contained unit of information transmitted as a whole, in a format required by DNS protocol. Similarly, the '040 patent provides an example in which the service request packet is an HTTP GET message—a self-

contained unit of information transmitted as a whole according to HTTP. Ex. 1001, FIG. 2 (showing the format of the HTTP service request message).

160. Accordingly, in my opinion, Treuhaft discloses, or at least suggests, "receiving a service request packet sent by a terminal device, wherein the service request packet carries … a server domain name indicating the service server requested by the service request packet sent by the terminal device," as claimed.

161. In my opinion, Treuhaft further discloses or renders obvious that the modified DNS query (claimed service request packet) also carries a terminal domain name indicating the terminal device (host device 105), as claimed. Specifically, before the host device 105 sends the DNS query to name server 102, "[i]n step 520, the DNS query is modified with control information." Ex. 1008, ¶ [0058]; *see also id.*, ¶ [0067] ("control information may be encoded into an individual DNS query that enables a DNS nameserver to identify DNS resolution options, filters, or features to apply when resolving the individual DNS query"). In my opinion, this control information included in the modified DNS query discloses, or at least suggests, the claimed terminal domain name.

162. Under the broadest reasonable interpretation, "terminal domain name indicating the terminal device" includes an identifier associated with an owner of the terminal device. *See* Ex. Ex. 1012, 4 (definition of "domain name"); Ex. 1001, 3: 32-51 (describing a terminal domain name as a "unique identifier" of the terminal device). Treuhaft discloses several examples in which the control information serves as an identifier associated with an owner of the terminal device: "[t]he control information may specify … a user or subscriber identifier, a device identifier, or the like." Ex. 1008, ¶ [0036]. In my opinion, each of the user identifier, subscriber identifier, or device identifier serves as an identifier of the owner of the terminal device—in Treuhaft's case, host device 105—and thus discloses or suggests the claimed terminal domain name. Indeed, a domain name is "like" a user identifier, subscriber identifier, or device identifier in that it is a name or label of the device with which it is associated and/or its owner. Thus, in my opinion, Treuhaft at least contemplates a domain name in its description of the control information.

163. Treuhaft uses "user or subscriber" interchangeably to refer to "a user or subscriber of the OpenDNS service [who] set[s] one or more preferences or selections for how the options are to be enabled or otherwise applied when DNS nameserver 120 resolves DNS queries associated with the user," *Id.*, ¶ [0028]; *see also id.*, ¶¶ [0008], [0024], [0027], [0029], [0036], [0060]. And

when the host device 105 attempts to access a server at a certain domain name, the DNS name server 120 applies the corresponding user's preferences or selections when resolving the DNS query for the address of that domain name in step 540 of Figure 5B. *See id.*, ¶¶ [0065]; *see also id.*, ¶¶ [0028], [0035], [0039]. Thus, in my opinion, Treuhaft's user or subscriber information, included in the modified DNS query as control information, discloses the claimed terminal domain name because it identifies an owner associated with the host device 105 who is controlling the host device 105's access to the Internet.

164. Similarly, in my opinion, the device identifier of Treuhaft alternatively or additionally discloses or suggests the claimed terminal domain name. As highlighted in Figure 4 below, Treuhaft discloses the modified DNS query 400 may contain a device ID 480 "provided in the additional section of [the modified] DNS query." *Id.*, ¶ [0053].

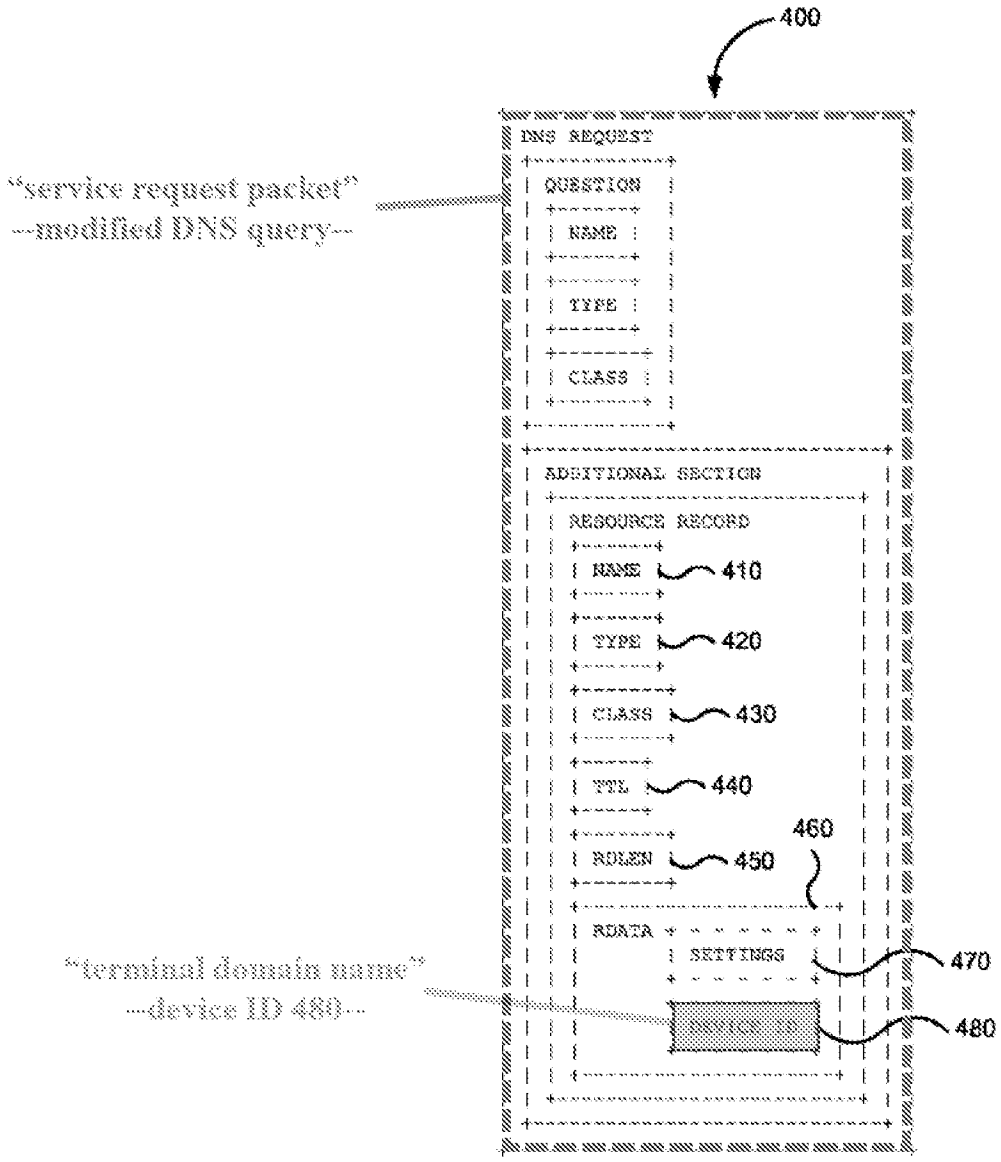

Treuhaft, FIG. 4*

In my opinion, the device ID 480 likewise corresponds to an owner associated with the host device 105 because "host device 105 can supply a device ID to DNS nameserver 120 by including DEVICE ID 480." *Id.* That is, the host device supplies its own device ID 480 when forming the modified DNS query 400. And, as explained above, the host device 105 seeking access to a certain domain name has an associated owner or subscriber to the DNS service who has supplied "one or more preferences or selections for how the options are to be enabled or otherwise applied when DNS nameserver 120 resolves DNS queries associated with the user." *Id.*, ¶ [0028]; *see also id.*, ¶¶ [0008], [0024], [0027], [0029], [0036], [0060]. Accordingly, Treuhaft's device ID 480, included in the modified DNS query as control information, also discloses or suggests the claimed terminal

domain name because it identifies an owner associated with the host device 105 who is controlling the host device 105's access to the Internet.

165. In another example disclosing or suggesting claimed terminal domain name, highlighted in Figure 4 below, Treuhaft teaches that the modified DNS query may contain a "CLIENTID for control of user, device, or vendor-specific DNS server behavior" within "the additional data section of a request." *Id.*, ¶ [0042]; *see also id.*, ¶¶ [0043]-[0045].

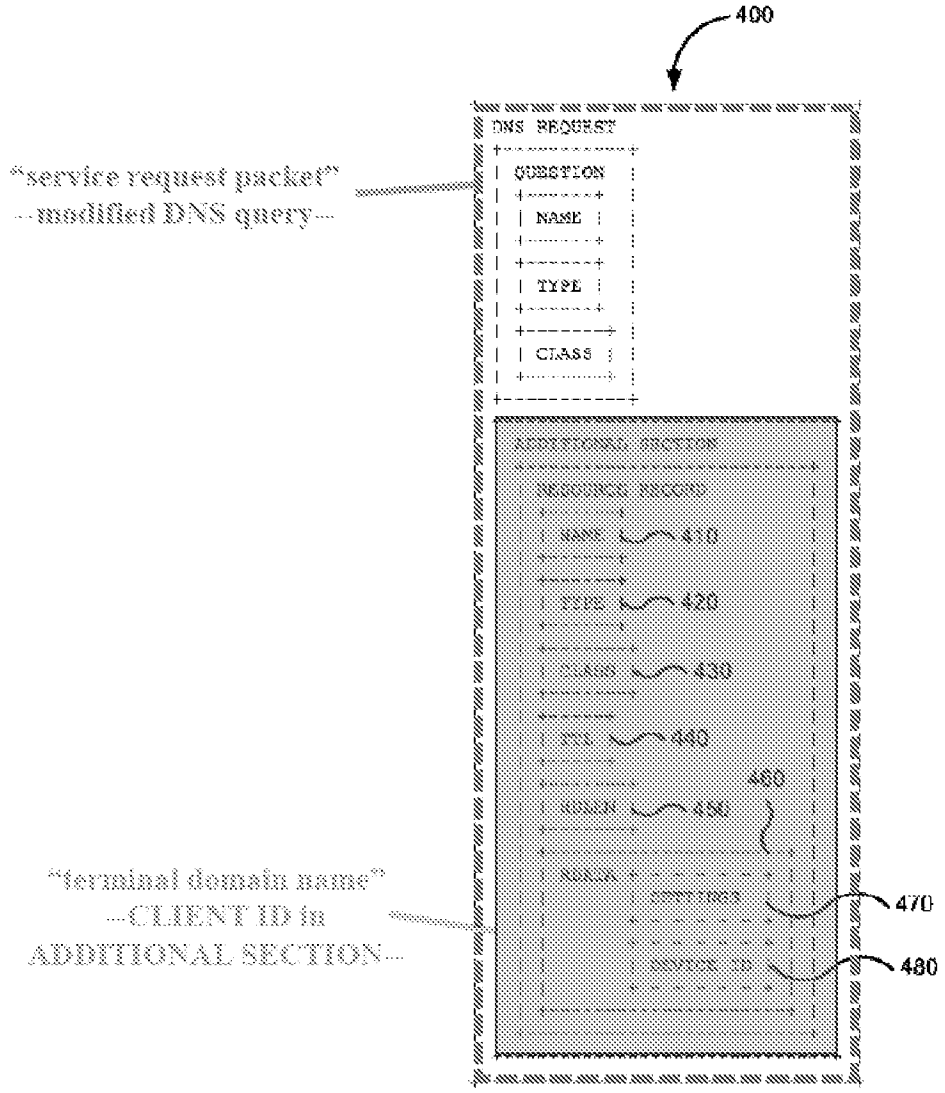

Treuhaft, FIG. 4*

Similar to the other control information discussed above, Treuhaft also uses the CLIENTID to identify the user or subscriber associated with the host device 105 and control the host device's access to the Internet accordingly. *See* Ex. 1008, ¶¶ [0042]-[0045].

166. Accordingly, as outlined above, the modified DNS query of Treuhaft discloses the claimed service request packet as follows:

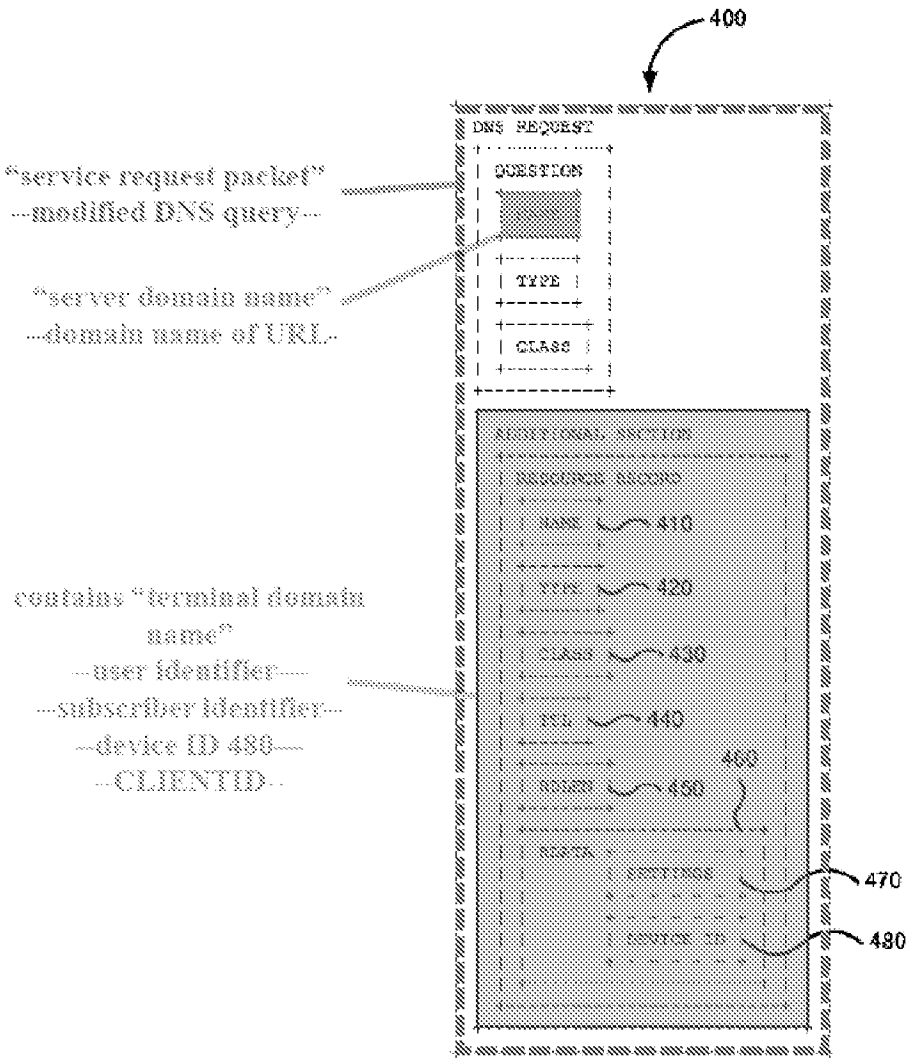

Treuhaft, FIG. 4*

167. To the extent it is argued Treuhaft's control information (e.g., user identifier, subscriber identifier, device ID, or CLIENTID) does not disclose a terminal domain name in the narrow sense of a hierarchical domain name in the form server.organization.type, *see* Ex. 1012, 4 (definition of domain name), in my opinion, it would have been obvious to a POSA to modify Treuhaft to use such a hierarchical domain name of the host device 105 as the control information. Treuhaft teaches that the purpose of the host device identifier contained in the DNS query is to "enable[] the domain name service to retrieve subscriber information" which "include[s] preferences or other settings for how a user or subscriber wishes to control domain name resolution

within the DNS resolution features." *Id.*, ¶ [0060]; *see also id.*, ¶¶ [0039] ("host device 105 may encode within a DNS query an identifier, such as an account ID or index, that specifies where DNS nameserver 120 can find the preferences or subscriber information used by options 27").

168. Thus, the purpose Treuhaft's control information is to identify the owner or subscriber associated with the host device 105 seeking access to the Internet. In my opinion, a POSA would have understood that any piece of information associated with the user or subscriber could be used as the control information in the DNS query, such as an IP address or hierarchical domain name. Indeed, according to its dictionary definition, a hierarchical domain name serves the exact purpose of Treuhaft's user identifier, subscriber identifier, device identifier, or other control information: "An address of a network connection that identifies the owner of that address in a hierarchical format." Ex. 1012, 4. And hierarchical domain names were routinely used for this purpose long before the '040 Patent. In my opinion, a POSA would have found it obvious to use a hierarchical domain name for the host device 105 as an alternative, or in addition to, Treuhaft's user identifier, subscriber identifier, device identifier, or other type of control information.

169. Accordingly, in my opinion, Treuhaft discloses, or at least suggests, "receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device … ," as claimed.

c. [1.2] *"resolving the received server domain name to obtain a service server Internet protocol (IP) address; and"*

170. In my opinion, Treuhaft discloses, or at least suggests, this element. Specifically, Treuhaft discloses resolving the received server domain name (domain name associated with URL contained in the modified DNS query) to obtain a service server Internet protocol (IP) address. Treuhaft explains, "in step 530, the DNS query is received. In step 535, the DNS query is parsed or otherwise processed at the DNS nameserver to determine the domain name and the control information." Ex. 1008, ¶ [0063]. Then, "[i]n step 540, an IP address is determined using one or more DNS resolution options or features and the control information. In one example, the domain name is resolved to its corresponding IP address." *Id.*, ¶ [0064].

171. Accordingly, Treuhaft discloses, or at least suggests, "resolving the received server domain name to obtain a service server Internet protocol (IP) address," as claimed.

d. [1.3] *"discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list,"*

172. In my opinion, Treuhaft discloses, or at least suggests, this element. As explained below, the DNS name server 120 of Treuhaft maintains subscriber information 208 for various users or subscribers, which corresponds to the claimed "preset list." *See id.*, ¶¶ [0028], [0029], [0034], [0036], [0039], [0054], [0060], [0064], FIG. 2 (subscriber information 280). If the subscriber information 208 for the user or subscriber indicates to block access to the IP address resolved from the modified DNS query—such as an IP address for an inappropriate website—the DNS name server 120 of Treuhaft discards the modified DNS query by not returning that IP address to the host device 105. *See, e.g., id.*, ¶¶ [0027], [0028].

173. Specifically, Figure 2 of Treuhaft shows the DNS name server 120 storing subscriber information 280 for the users or subscribers of the system in memory 220:

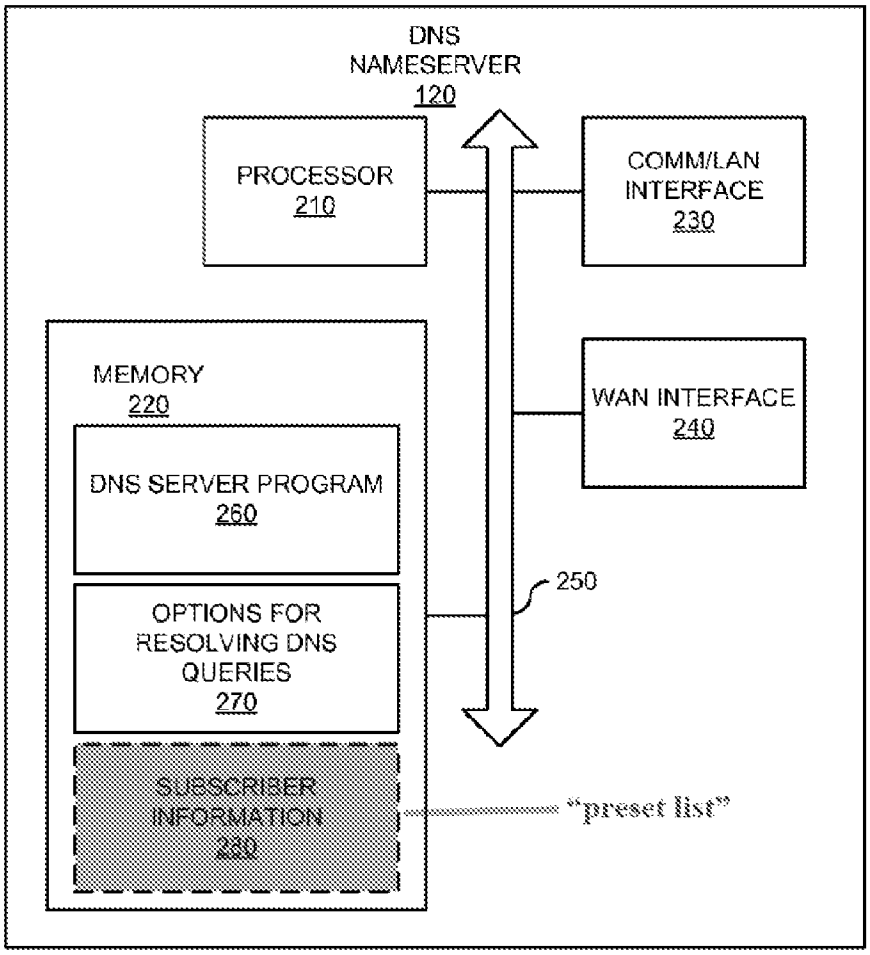

FIG. 2

Treuhaft, FIG. 2*

As highlighted above, the subscriber information 280 corresponds to the claimed preset list. The subscriber information can include preferences or other settings for how a user or subscriber wishes to control domain name resolution within the DNS resolution features." *Id.*, ¶ [0060]. "For example, a user or subscriber may establish subscriber information that instructs DNS nameserver 120 to alter responses to DNS requests that are associated with adult web sites, potential phishing or pharming sites, and other sites deemed inappropriate by the user or containing material illegal in the country of the user." *Id.*, ¶ [0028]. The "subscriber information associated with the user may be used to alter the IP address in a DNS response that the user receives." *Id.*

174. In steps 535-550 of Figure 5, the name server 120 of Treuhaft applies the subscriber information for the user or subscriber to the modified DNS query in determining how to respond

to the modified DNS query. *See id.*, ¶¶ [0064]-[0066]. In my opinion, this process corresponds to the claimed "discarding the service request packet":

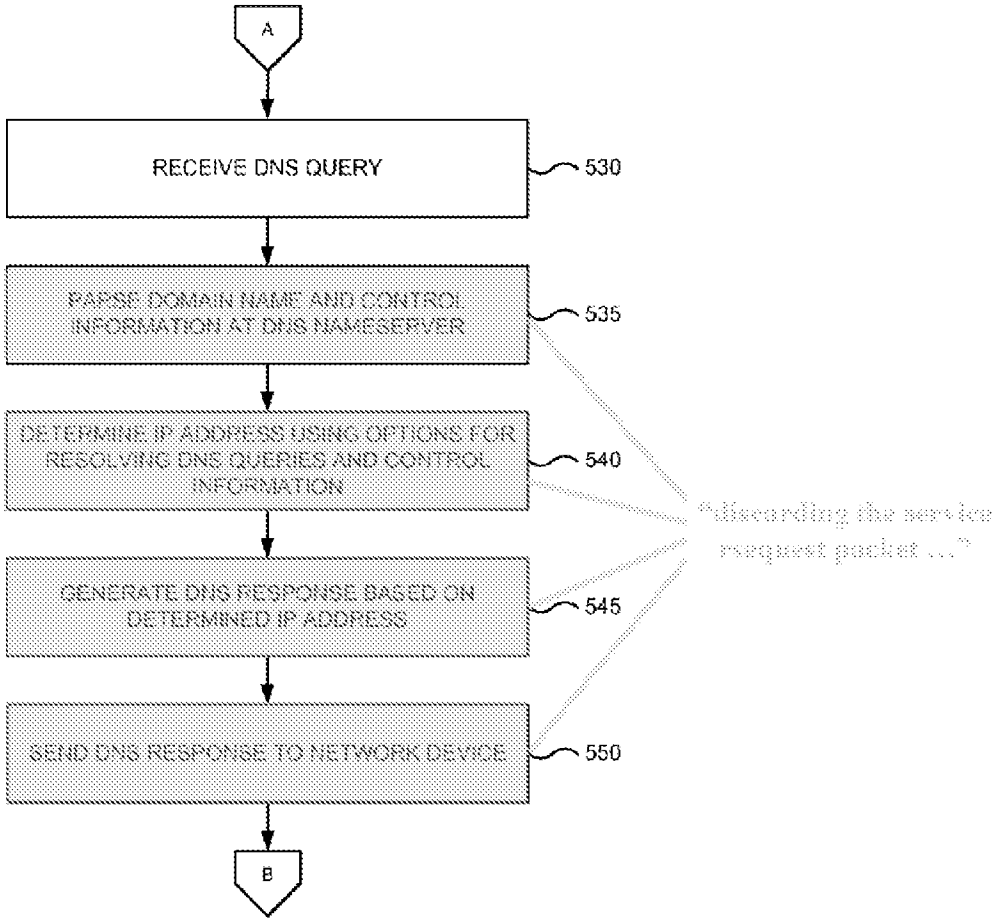

FIG. 5B

Treuhaft, FIG. 5B*

175. Specifically, after receiving the modified DNS query, the DNS name server 120 parses the control information in the modified DNS query to identify the particular user or subscriber that sent the DNS query, and retrieves that user or subscriber's subscriber information 280 in step 535. *Id.*, ¶ [0064]. Then, in step 540, using the user or subscriber's subscription information 280, the DNS name server 120 "make[s] a decision whether to use the corresponding IP address or another IP address when generating a DNS response based on applying one or more DNS resolution options or features". *Id.* ¶ [0065]. For example, rather than return the resolved IP address requested by the host device 105, "DNS nameserver 120 may determine to substitute the

IP address of a website that provides information why the domain name is being block[ed], forwarded, filtered, or otherwise includes material the user has expressed a desire to control." *Id.*, ¶ [0065]. Then, in steps 545 and 550, the DNS name server 120 respectively generates a DNS response "substitut[ing] [the] IP address based on applying one or more of the available DNS resolution options, filters, or features" and sends the DNS response with the substituted IP address to the host device 105. *Id.*, ¶ [0066].

176. As explained above, the broadest reasonable interpretation of "discarding the service request packet" includes preventing unauthorized access to the resolved service server IP address. *Supra* Section VI.A. Treuhaft's process in steps 535-550 of Figure 5B prevents the host device 105 from unauthorized access to the resolved IP address. This is because the DNS name server 120 alters or substitutes the resolved IP address for a different IP address in the DNS response if the subscriber information deems access to that IP address unauthorized. Ex. 1008, ¶¶ [0065], [0066]; *see also id.*, ¶¶ [0028], [0035] For example, Treuhaft may instead "substitute the IP address of a website that provides information why the domain name is being block[ed], forwarded, filtered, or otherwise includes material the user has expressed a desire to control." Ex. 1008, ¶ [0065]; *see also id.*, ¶ [0032] (the DNS name server 120 "respond[s] with another IP address that, for example, redirects the user to a website with additional information for the reason why the corresponding IP address was not returned"). By responding to the DNS query with a different IP address than the one requested and blocking the requested IP address, Treuhaft prevents unauthorized access to the resolved IP address. Thus, in my opinion, Treuhaft discloses or suggests "discarding the service request packet," as claimed.

177. Moreover, even if the claimed discarding were construed narrowly to mean providing no DNS response at all to the DNS query, this option is suggested based on Treuhaft's disclosure of "blocking" or "not returning" the IP address for an unauthorized site. *See id.*, ¶¶ [0032], [0065].

178. In my opinion, Treuhaft's subscriber information discloses, or at least suggests, a "preset list" containing "a preset service server IP address corresponding to the received terminal domain name," as claimed. In Treuhaft, the DNS name server 120 applies the subscriber information to "make a decision whether to use the corresponding [resolved] IP address or another IP address" if the resolved IP address is not authorized. *Id.*, ¶ [0065]; *see also id.* ¶¶ [0065], [0066]

(if the resolved IP address is not authorized, the name server 120 returns a "substitute IP address"). Treuhaft gives an example in which, applying the subscriber information, the "DNS nameserver 120 may respond with the [resolved] IP address of 'www.cnet.com' .... or may respond with another IP address that, for example, redirects the user to a website with additional information for the reason why the corresponding IP address was not returned." *Id.*, ¶ [0032].

179. Based on this disclosure, in my opinion, a POSA would have understood that the DNS name server 120 of Treuhaft checks the resolved IP address for "www.cnet.com" against known authorized/unauthorized IP addresses in deciding whether to return the resolved IP address for "www.cnet.com" or another IP address. That is, if the resolved IP address is authorized for the user/subscriber of the host device 105—i.e., "belongs to a preset service server IP address corresponding to the received terminal domain name," as claimed—the DNS name server 120 returns the resolved IP address. *Id.* Otherwise, the DNS name server 120 discards the DNS query by returning a different IP address, thus preventing access to the resolved IP address as discussed above.

180. Given Treuhaft's IP address comparison, Treuhaft discloses or at least suggests that the subscriber information includes a "preset list" of one or more authorized or unauthorized IP addresses, as claimed. Indeed, in order to make this IP address comparison, Treuhaft necessarily must store a list of authorized/unauthorized IP addresses somewhere, and the subscriber information 208 for the particular user/subscriber is the most logical place. As explained above, the DNS name server 120 stores the subscriber information 280 in memory 220. *Id.*, ¶ [0034], FIG. 2. In my opinion, it would have been obvious to implement Treuhaft's subscriber information as a preset list of authorized/unauthorized IP addresses. Moreover, it would have been obvious to store this list of authorized/unauthorized IP addresses in the memory 220 as part of the subscriber information 280 for the particular user/subscriber.

181. If it is argued that Treuhaft's "preset list" functions as a blacklist for discarding the DNS query if the resolved IP address belongs to the list, rather than "does not belong" as claimed, it would have been obvious to alternatively or additionally implement Treuhaft's list as a whitelist. Long before the claimed priority date of the '040 patent, blacklists and whitelists were known and used interchangeably and/or together in the same system to control access to sites. Depending on the knowledge of the administrator configuring the access controls for users or subscribers of the

system, it would have been obvious to use a whitelist of authorized IP addresses, a blacklist of unauthorized IP addresses, or a combination of both a whitelist and a blacklist. For example, if the administrator wanted to configure Treuhaft's system so that a user or subscriber may only access certain known IP addresses, the administrator would have found it obvious to use a whitelist of authorized IP addresses. But if the administrator wanted to allow access to all sites except those specifically deemed inappropriate, for example, the administrator would have found it obvious to use a blacklist.

182. Accordingly, for at least the reasons above, Treuhaft discloses or suggests "discarding the service request packet if the resolved service server IP address does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list," as claimed.

e. [1.4] *"wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device."*

183. In my opinion, Treuhaft discloses, or at least suggests, this element. As explained above, the subscriber information 280 of Treuhaft discloses or suggests a preset list of authorized/unauthorized IP addresses that a given user or subscriber has or does not have authorization to access. *Supra* Section X.B.1.d. Moreover, in my opinion, Treuhaft's list of authorized/unauthorized IP addresses for users or subscribers discloses or suggests the claimed "corresponding[] … plurality of accessible service server IP addresses under an access authority of the terminal device," as recited in element [1.4].

184. Treuhaft explains that its system allows "DNS resolution [to] be controlled on a per-request basis for each individual user or device." Ex. 1008, ¶¶ [0008], [0024]. This suggests the DNS name server 120 contains subscriber information 280 for multiple users or subscribers, as the DNS name server 120 would need this information to process DNS queries on a user-by-user or subscriber-by-subscriber basis. As discussed above, the subscriber information 280 includes, or suggests, a list of authorized/unauthorized IP addresses for its users or subscribers. Thus, it would have been obvious to arrange Treuhaft's list such that the user identifier, subscriber identifier, device ID, CLIENTID, or hierarchical domain name (claimed terminal domain name) associated with each host device 105 (claimed terminal device) maps to the corresponding

authorized/unauthorized IP address(es) for that user or subscriber. In my opinion, mapping the host device identifiers to their corresponding authorized/unauthorized IP addresses in this manner would allow the DNS name server 120 to check the resolved IP address against the authorized/unauthorized IP addresses for that user or subscriber when processing a DNS query from that user or subscriber.

185. Accordingly, Treuhaft discloses, or at least suggests, "wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with a plurality of accessible service server IP addresses under an access authority of the terminal device," as claimed.

2. Dependent Claim 4

a. [4.1] ***"The method according to claim 1, wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:***

if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device."

186. Claim 4 depends on independent claim 1 and recites the additional elements quoted above. Treuhaft anticipates and renders obvious independent claim 1 for the reasons discussed above. *Supra* Section X.B.1. Moreover, in my opinion, Treuhaft discloses, or at least suggests, the additional elements of claim 4.

187. As explained above, Treuhaft discloses or suggests that the DNS name server 120 uses a list of authorized/unauthorized IP addresses for each user or subscriber of the system to determine whether the resolved IP address is authorized/unauthorized. *Supra* Sections X.B.1.d, e. If the resolved IP address authorized, the "DNS nameserver 120 … use[s] the corresponding IP address of the domain name"—i.e., the resolved IP address—in the DNS response to the host device 105 in steps 545 and 550. Ex. 1008, ¶ [0066]. As shown in Figure 5C of Treuhaft, the host device 105 subsequently receives the DNS response from the name server 120 (step 555), determines the resolved IP address from the DNS response (step 560), and forwards the IP address to an application running on the host device 105 (step 565). *See id.*, ¶ [0067], FIG. 5C.

188. In my opinion, a POSA would have understood that, upon forwarding the IP address to the application on the host device 105, the application connects to the server associated with that IP address. Indeed, Treuhaft teaches that the application running on the host device, which sent the DNS query for the IP address, is a web browser. *See id.*,¶¶ [0025], [0057]. And a web browser is an application, used to browse the Internet, that sends a DNS query for an IP address to a DNS server, receives a DNS response containing the IP address from the DNS server, and connects to a server located at the IP address. If it is argued that Treuhaft does not expressly disclose connecting to the server using the IP address, a POSA would have found it obvious to have the browser application connect to the server located at the returned IP address because Treuhaft seeks to enable Internet connections with DNS-based access controls, and the purpose of a browser is to access the Internet using IP addresses returned from DNS servers. *See, e.g.*, *id.*, ¶¶ [0022]-[0024]

189. Accordingly, in my opinion, any combination of steps 545, 550 (Figure 5B, reproduced below) performed by the name server 120 and steps 555-565 (Figure 5C) performed by the host device 105 discloses or suggests "if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection … ," as recited in claim 4:

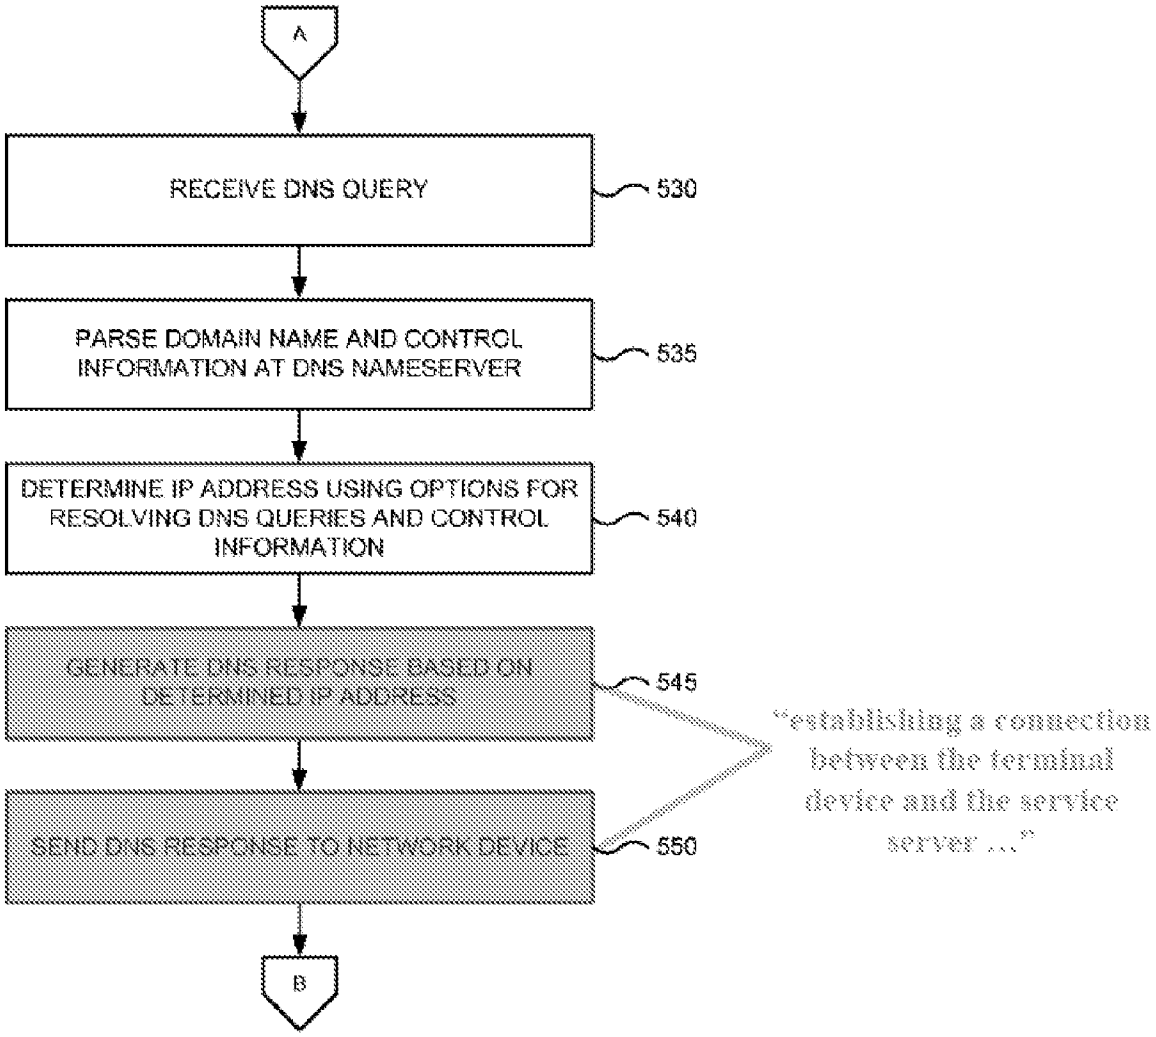

FIG. 5B

Treuhaft, FIG. 5B*

190. In my opinion, under the broadest reasonable interpretation standard, the DNS name server 120's sending the DNS response with the resolved IP address to the host device 105 facilitates "establishing a connection," as claimed. This is because, as explained above, the application running on the host device 105 connects to the intended server using the IP address contained in the DNS response upon receiving the DNS response. *See, e.g.*, Ex. 1008, ¶ [0067], FIG. 5C; *see also id.*, ¶¶ [0025], [0057]. That is, the process of returning the DNS response to the host device enables the establishment of a connection between the host device 104 (claimed terminal device) and the server located at the IP address (claimed service server corresponding to the service server IP address).

191. Accordingly, in my opinion, Treuhaft discloses or suggests "wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises: if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device," as recited in claim 4.

3. Dependent Claim 5

a. [5.1] ***"The method according to claim 1, wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:***

if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device."

192. Claim 5 depends on independent claim 1 and recites the additional elements quoted above. Treuhaft anticipates and renders obvious independent claim 1 for the reasons discussed above. *Supra* Section X.B.1. Moreover, in my opinion, Treuhaft discloses, or at least suggests, the additional elements of claim 5 because, as explained below, the DNS name server 120 uses a resource record contained in the DNS query to determine the DNS extension capabilities of the host device 105 (claiming determining a service type).

193. As Treuhaft explains, the DNS query 400 contains fields the host device 105 uses "to advertise its own extended capabilities to the message receiver (e.g., DNS nameserver 120)." *See* Ex. 1008, ¶¶ [0040]-[0041], FIG. 4. "This may be accomplished through the inclusion of an OPT pseudo-RR in the additional data section of a request or response. The OPT pseudo-RR may include one or more EDNS options." *Id.*, ¶ [0041]. OPT pseudo-RR refers to an options resource record contained in a DNS query sent according to the Extension Mechanisms for DNS (EDNS) specification. A device uses this resource record in a DNS query to identify DNS extension capabilities to the DNS server. For example, one such extension mechanism is DNS Security Extensions (DNSSEC), securing data exchanged over DNS using cryptography.

194. In my opinion, a POSA would have understood that the host device 105 in Treuhaft uses the options resource record to indicate to the DNS name server 120 that it supports DNSSEC or another type of DNS extension (claimed service type). *See id.*, ¶ [0041]. In the case of a DNSSEC-enabled host device 150, for example, the host device 105 includes information in the options resource record of the DNS query indicating that it supports DNSSEC. Upon receiving the DNS query, the DNS name server 120 unpacks the DNS query, and determines from this resource record that the host device 105 supports DNSSEC. Accordingly, the DNS server responds to the host device 105 according to the DNSSEC protocol, such as by engaging in a cryptographic handshake routine with the host device 105 and/or encrypting its DNS response. Thus, according to EDNS protocol, the DNS name server 120 of Treuhaft "determin[es] a service type of the service request," as claimed, based on information contained in the options resource record. *See id.*, ¶¶ [0040]-[0041].

195. In my opinion, Treuhaft further discloses that the DNS name server 120 determines the service type of the DNS query "according to the terminal domain name of the terminal device," as claimed. For example, Treuhaft explains that the host device 105 may include its CLIENTID in the options resource record to identify its DNS extension capabilities to the DNS name server 120:

> In some embodiments, host device 105 can define a new EDNS option called CLIENTID for control of user, device, or vendor-specific DNS server behavior. The CLIENTID option may appear in an OPT pseudo-RR in the additional data section of a request. In general, a CLIENTID option applies to the DNS request that it accompanies. Thus, the CLIENTID can allow a per-request control of each DNS message.

Id., ¶ [0042]. Thus, the DNS name server 120 of Treuhaft uses the CLIENTID in the options resource record of the DNS query to determine the host device 105's (claimed terminal device) DNS extension capabilities (claimed service type). *Id.* As explained above, the CLIENTID corresponds to the claimed terminal domain name because it identifies the user or subscriber associated with the host device 105. *Supra* Section X.B.1.b; *see also* Ex. 1008, ¶¶ [0042]-[0045]. And because the DNS name server 120 uses the CLIENTID (claimed terminal domain name) in the options resource record of the DNS query to determine the host device 105's DNS extension capabilities, it is my opinion that Treuhaft discloses or suggests determining the service type of the DNS query "according to the terminal domain name of the terminal device," as claimed.

196. For at least the reasons above, it is my opinion that Treuhaft discloses or suggests "wherein, after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises: if the resolved service server IP address belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device," as recited in claim 5.

4. Independent Claim 6

197. As set forth below, it is my opinion that Treuhaft discloses or suggests of the elements of independent claim 6 for reason similar to those discussed above for independent claim 1.

a. [6.pre] ***"A deep packet inspection (DPI) device comprising a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method comprising:"***

198. To the extent the preamble is limiting, in my opinion, Treuhaft discloses this element. Specifically, the DNS name server 120 of Treuhaft corresponds to the claimed DPI device:

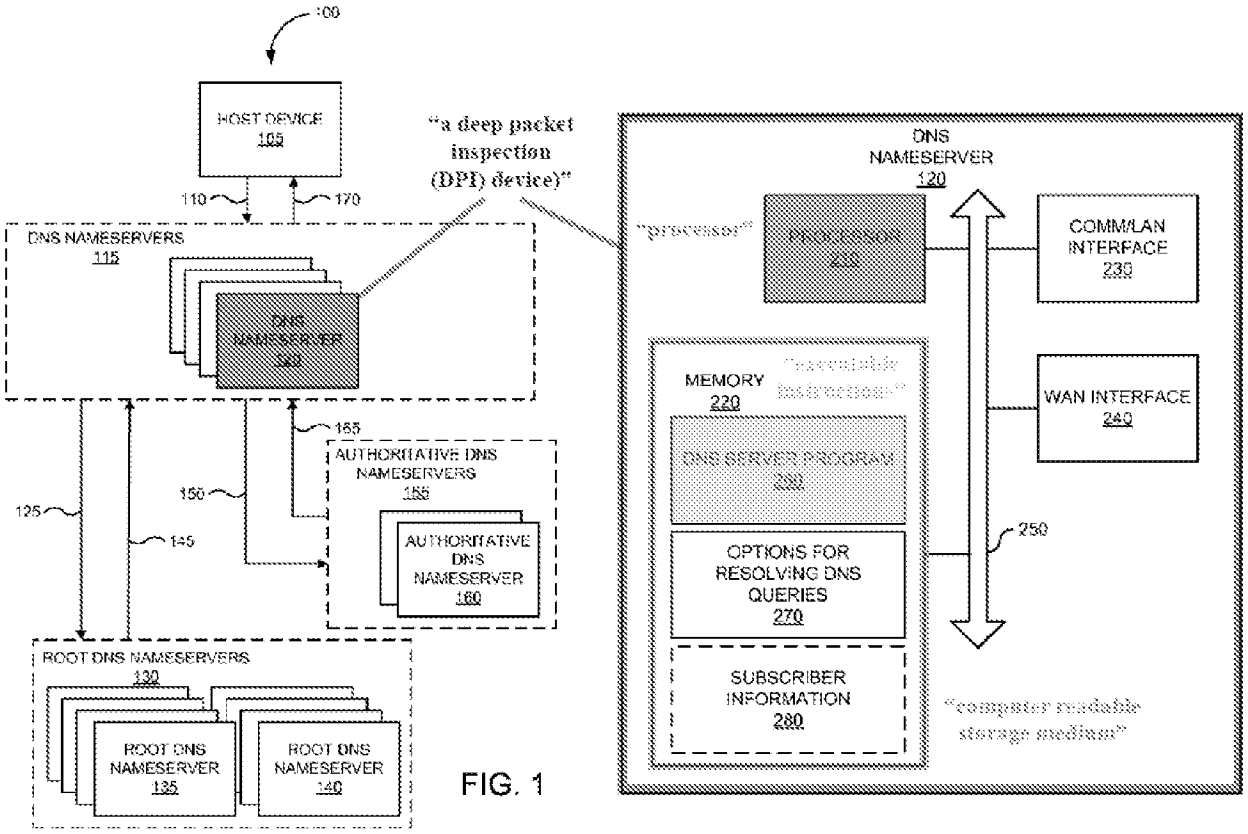

FIG. 1

FIG. 2

Treuhaft, FIGS. 1, 2*

As highlighted above, the DNS name server 120 has a processor 210 (claimed processor) and a memory 200 (claimed computer readable storage medium) storing a DNS server program 260 (claimed executable instructions), executed by the processor 210 to perform the functions of the DNS name server 120. Ex. 1008, ¶¶ [0033]-[0034], [0038].

199. Accordingly, in my opinion, Treuhaft discloses or suggests "[a] deep packet inspection (DPI) device comprising a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method," as claimed.

b. [6.1] *"receiving a service request packet sent by a terminal device, wherein the service request packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request packet sent by the terminal device;"*

200. *See* the discussion above for element [1.1]. *Supra* Section X.B.1.b.

c. [6.2] *"resolving the server domain name to obtain a service server Internet protocol (IP) address; and"*

201. *See* the discussion above for element [1.2]. *Supra* Section X.B.1.c.

d. [6.3] *"discarding the packet if the service server IP address resolved does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset list,"*

202. *See* the discussion above for element [1.3]. *Supra* Section X.B.1.d.

e. [6.4] *"wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with accessible service server IP addresses under an access authority of the terminal device."*

203. *See* the discussion above for element [1.4]. *Supra* Section X.B.1.e.

5. Dependent Claim 9

a. [9.1] *"The DPI device according to claim 6, wherein after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:*

if the service server IP address resolved belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, establishing a connection between the terminal device and the service server corresponding to the service server IP address, to enable the service server to provide a service corresponding to the service request of the terminal device to the terminal device."

204. Claim 9 depends on independent claim 6 and recites the additional elements quoted above. Treuhaft anticipates and renders obvious independent claim 6 for the reasons discussed above. *Supra* Section X.B.4. Moreover, claim 9 recites subject matter virtually identical to that discussed above for dependent claim 4. Accordingly, for the reasons discussed above in connection with claim 4, it is my opinion that Treuhaft discloses, or at least suggests, the subject matter of claim 9. *Supra* Section X.B.2. For example, as explained above, the DNS name server in Treuhaft facilitates or enables establishing the connection by returning the resolved IP address to the host device, which uses the returned IP address to connect to the server.

6. Dependent Claim 10

a. [10.1] ***"The DPI device according to claim 6, wherein after the resolving the received server domain name to obtain the service server Internet protocol (IP) address, the method further comprises:***

if the service server IP address resolved belongs to the preset service server IP address corresponding to the received terminal domain name in the preset list, determining a service type of the service request according to the terminal domain name of the terminal device."

205. Claim 10 depends on independent claim 6 and recites the additional elements quoted above. Treuhaft anticipates and renders obvious independent claim 6 for the reasons discussed above. *Supra* Section X.B.4. Moreover, claim 10 recites subject matter virtually identical to that discussed above for dependent claim 5. Accordingly, for the reasons discussed above in connection with claim 5, it is my opinion that Treuhaft discloses, or at least suggests, the subject matter of claim 10. *Supra* Section X.B.3.

7. Independent Claim 11

206. As set forth below, in my opinion, Treuhaft discloses or suggests of the elements of independent claim 11 for reasons similar to those discussed above for independent claims 1 and 11.

a. [11.pre] ***"A system, comprising: a deep packet inspection (DPI) device; and a terminal device:"***

207. As highlighted in Figure 1 below, Treuhaft discloses a DNS system 100 (claimed system) comprising a DNS name server 120 (claimed DPI device) and a host device 105 (claimed terminal device):

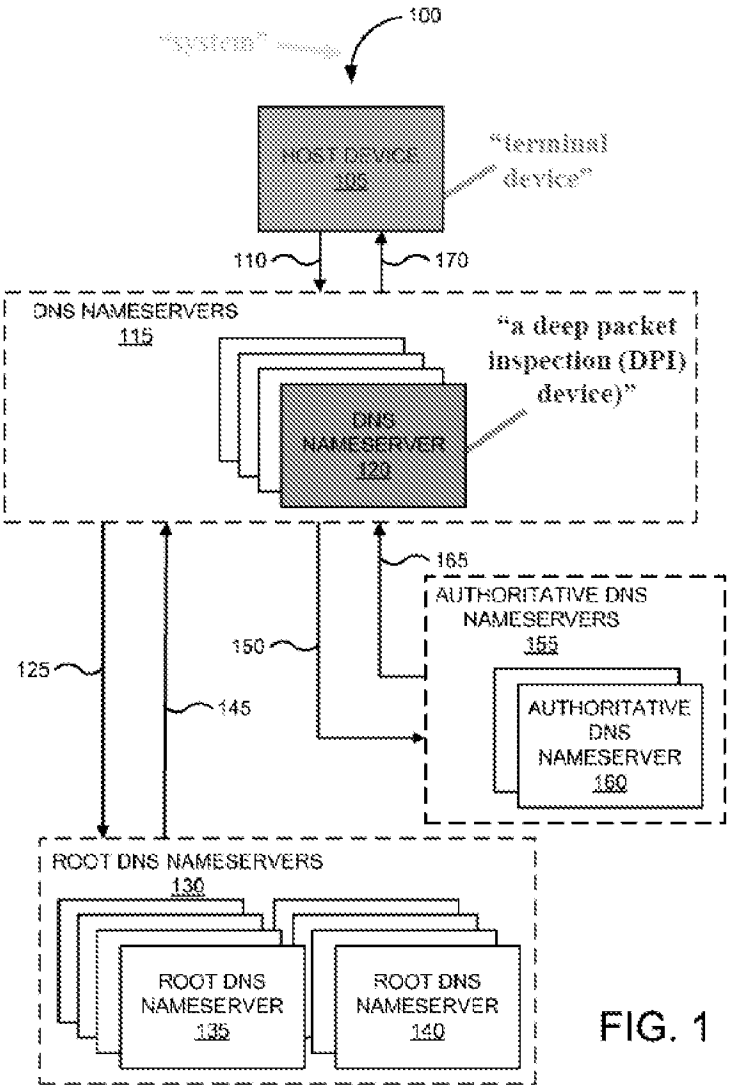

FIG. 1

Treuhaft, FIG. 1*

208. Accordingly, in my opinion, Treuhaft discloses or suggests "[a] system, comprising: a deep packet inspection (DPI) device; and a terminal device," as claimed.

b. [11.1] [*a terminal device*] "*configured to send a service request packet to the DPI device, wherein the packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request sent by the terminal device;*" and

"*the DPI device having a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method comprising: receiving the service request packet*

sent by the terminal device;"

209. As I explain above for element [1.1], in Treuhaft, the host device 105 (claimed terminal device) is configured to send a DNS query (claimed service request packet) to the DNS name server 120 (claimed DPI device). *Supra* Section X.B.1.b. Additionally, the DNS query (claimed service request packet) carries control information (claimed terminal domain name) indicating the host device 105 (claimed terminal device) and a domain name of a URL (claimed server domain name) indicating a server to which the host device 105 seeks to connect (claimed service server required by the service request sent by the terminal device). *Id.*

210. As explained above for element [6.pre], the DNS name server 120 (claimed DPI device) has a processor 210 (claimed processor) and a memory 200 (claimed computer readable storage medium) storing a DNS server program 260 (claimed executable instructions), executed by the processor 210 to perform the functions of the DNS name server 120. Ex. 1008, ¶¶ [0033]-[0034], [0038]. And, part of that method includes receiving the DNS query (claimed service request packet) sent by the host device 105 (claimed terminal device). *Supra* Section X.B.1.b; *see also* Ex. 1008, ¶¶ [0063] ("In step 525, the modified DNS query is sent to a DNS nameserver. For example, the modified DNS query may be sent to DNS nameserver 120."), ¶ [0064] ("in step 530, the DNS query is received"), FIG. 5B (step 550), FIG. 5C (step 555).

211. Accordingly, in my opinion, Treuhaft discloses or suggests "a terminal device … configured to send a service request packet to the DPI device, wherein the packet carries a terminal domain name indicating the terminal device and a server domain name indicating a service server required by the service request sent by the terminal device" and "the DPI device having a hardware processor and a non-transitory computer readable storage medium including executable instructions that, when executed by the processor perform a method comprising: receiving the service request packet sent by the terminal device," as claimed.

c. [11.2] *"resolving the server domain name to obtain a service server Internet protocol (IP) address; and"*

212. *See* the discussion above for element [1.2]. *Supra* Section X.B.1.c.

d. [11.3] *"discarding the packet if the service server IP address resolved does not belong to a preset service server IP address corresponding to the received terminal domain name in a preset*

list,"

213. *See* the discussion above for element [1.3]. *Supra* Section X.B.1.d.

e. [11.4] *"wherein in the preset list the terminal domain name of each terminal device is correspondingly provided with accessible service server IP addresses under an access authority of the terminal device."*

214. *See* the discussion above for element [6.4]. *Supra* Section X.B.1.e.

C. Ground 4: Treuhaft in View Sorenson Renders Obvious Claims 1, 4-6, and 9-11 of the '040 Patent

215. It might be argued Treuhaft's subscriber information 280 constitutes a preset list of authorized/unauthorized server domain names, rather than IP addresses. On this basis, it might be argued that Treuhaft discloses determining whether the server domain name is on the preset list—not the IP address resolved from that domain name—and so Treuhaft does not disclose "discarding the service request packet if the [resolved] service server IP address does not belong to a preset service server IP address ... in a preset list," as recited in elements [1.3], [6.3], and [11.3]. This argument would be incorrect at least because, as discussed above, Treuhaft at least suggests that the subscriber information 280 contains a list of authorized and/or unauthorized IP addresses for each user or subscriber. *Supra* Sections X.B.1.d, e. In the case of an authorized IP address list, in my opinion, a POSA would have understood Treuhaft to disclose discarding the DNS query if the resolved IP address is not on the authorized list. And, moreover, a POSA would have found it obvious to use a list of authorized IP addresses, unauthorized IP addresses, or both for the purpose of controlling access to the Internet.

216. Nonetheless, even if such an argument were accepted, in my opinion, the additional disclosure of Sorenson renders obvious elements [1.3], [6.3], and [11.3]. Accordingly, as explained below, it is my opinion that the combination of Treuhaft in view of Sorenson further renders obvious claims 1, 4-6, and 9-11 of the '040 Patent.

1. Sorenson Discloses Discarding the Service Request Packet "if the [Resolved] Service Server IP Address Does Not Belong to a Preset Service Server IP Address ··· in a Preset List" (Elements [1.3], [6.3],

and [11.3])

217. In my opinion, Sorenson discloses or suggests elements [1.3], [6.3], and [11.3] by denying a connection if the IP address resolved from a domain name in the connection request is found on an IP address blacklist.

218. Specifically, Sorenson discloses a "system and method for blocking access by a network device to specific network resources by comparing a specific resource identifier against entries in a blacklist and facilitating a connection accordingly." Ex. 1009, Abstract. In Sorenson, the system receives "a call request for the establishment of a communication session between IP device 12 and associated service 20. *Id.*, ¶¶ [0027]; *see also id.*, ¶¶ [0031], [0032]. The IP device 12, call request, and service 20 of Sorenson respectively correspond to the claimed terminal device, service request packet, and service server. Like the service request packet of the '040 Patent, Sorenson's call request "include[s] a specific identifier such as an entered IP address, domain name, or conventional phone number or name resolved into one of an IP address or domain name," which corresponds to the claimed server domain name. *Id.*, ¶ [0031].

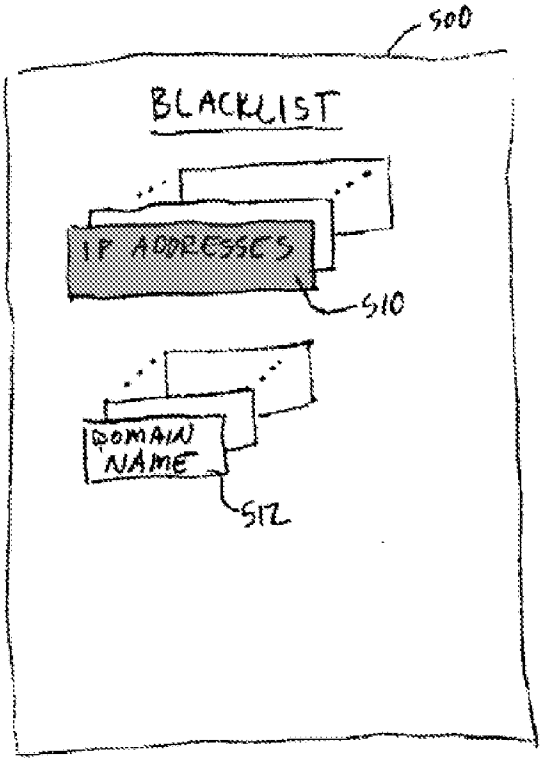

219. Before granting the call request, the system of Sorenson performs a two-stage blacklist check to determine whether to establish the connection or discard the request. *See id.*, ¶¶ [0031]-[0032], FIG. 6. As shown in Figure 2 of Sorenson above, the blacklist 500 contains both blacklisted domain name names 512 and blacklisted IP addresses 510. *See id.*, ¶ [0028], FIG. 2. In my opinion, the blacklist 500 of blacklisted IP addresses 510 corresponds to the claimed preset list.

220. First, as shown in step 610 of Figure 6, Sorenson performs a domain-name-blacklist check by "compar[ing] 610 the domain name against the blacklist 500´ (FIG. 2) to determine 612 if the domain name is located within the blacklist 500´." *Id.*, ¶ [0031]. "If the domain name utilized for initiating the call is located with the blacklist 500´, then the IP device denies 618 the completion of the call and may alternatively notify the user of such denial." But "[i]f the domain name is not on the blacklist, then" Sorenson performs an IP-address-blacklist check before establishing the connection. *Id.* Specifically, Sorenson "resolves … the domain name into an IP address for further comparison" in step 614, *id.*, and then "compares … the IP address against the blacklist 500'" in step 616, *id.*, ¶ [0032]. If the IP address is located within the blacklist 500', Sorenson denies the connection. *Id.*, ¶ [0032], FIG. 6 (step 618). But if the IP address is not found on the blacklist 500', Sorenson establishes the connection. *Id.*, ¶ [0032], FIG. 6 (step 622).

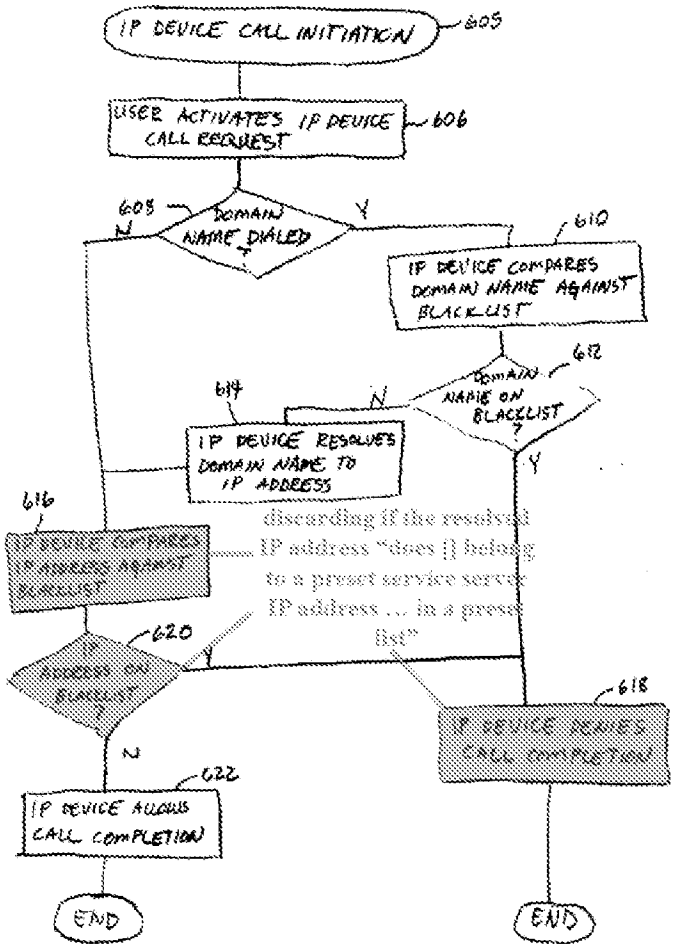

221. As highlighted in Figure 6 of Sorenson on the right, the combination of steps 616, 620, and 618—in which Sorenson determines whether the resolved IP address is on the IP backlist and denies the connection if it is—corresponds to "discarding the service request packet if the [resolved] service server IP address does not belong to a preset service server IP address … in a preset list," as recited in elements [1.3], [6.3], and [11.3].

2. Rationale to Combine Sorenson with Treuhaft

222. In my opinion, a POSA would have found it obvious and been motivated to use a list of authorized/unauthorized IP addresses in Treuhaft based on Sorenson's disclosure for several

reasons. The combination merely involves the use of known technique (Sorenson's IP address check) to improve similar devices (Treuhaft's system) in the same way and/or applying a known technique (Sorenson's IP address check) to a known system (Treuhaft) ready for improvement to yield predictable results.

223. Treuhaft and Sorenson are similar in several ways. For example, Treuhaft aims to deny DNS queries seeking access to sites "categorized as an adult web site, a potential phishing or pharming site, and a website whose content has been deemed inappropriate by the user or containing material illegal in the country of the user." Ex. 1008, ¶ [0027]; *see also id.* ¶¶ [0006], [0028]. Similarly, Sorenson seeks "to prevent access by user 14 to unauthorized or blacklisted services" on the Internet. Ex. 1009, ¶ [0023]. Thus, Treuhaft and Sorenson are analogous references having the same purpose to prevent users from accessing inappropriate sites, services, or other online resources.

224. Treuhaft and Sorenson also have similar structure and operation. For example, Treuhaft's DNS name server 120 stores subscriber information 280 used to control access to authorized/unauthorized sites for each user or subscriber of the system. *See, e.g.*, Ex. 1008, ¶¶ [0023], [0028], [0029], [0034], [0036], [0039], [0064], [0065]. Similarly, Sorenson's system maintains an IP address 510 and domain name 512 blacklist 500 used to control user access to online resources. *See* Ex. 1009, ¶¶ [0023], [0025], [0026], [0028]-[0032]. Both systems also receive and screen connection requests before or allowing or denying them. *Compare* Ex. 1008, ¶¶ [0064]-[0067], FIG. 5B *with* Ex. 1009, ¶¶ [0031]-[0032], FIG. 6. Accordingly, Treuhaft and Sorenson describe structurally- and functionally- similar systems designed to achieve a similar purpose.

225. To the extent it is argued that Treuhaft only checks the domain name before allowing or denying the connection, Sorenson improves upon Treuhaft by using a blacklist to check both the domain name and the IP addressed resolved from that domain name before allowing or denying a connection. Ex. 1009, ¶¶ [0027], [0031], [0032]. In my opinion, a POSA have understoods that domain names and IP addresses do not necessarily have a one-to-one correspondence, as multiple domain names might resolve to the same IP address, and a given domain may resolve to multiple IP addresses. In some cases, for example, the same inappropriate website may have multiple domain names, or different DNS servers may resolve different domain

names to that same inappropriate address. An administrator might be aware of some of those domain names but not others, and thus only include the known domain names in a domain name blacklist. If a user later issues a connection request using one of the unknown domain names for the inappropriate site not on the blacklist, the system would allow the connection to the inappropriate site. Sorenson's technique of checking the resolved IP address as well as the domain name, however, would catch and deny such inappropriate connection requests that Treuhaft's system might otherwise allow. Accordingly, Sorenson's application of both a domain name and an IP address blacklist improves upon Treuhaft, to the extent Treuhaft only discloses checking the domain name before allowing or denying a connection. Thus, in my opinion, a POSA would have had motivation to make the combination.

226. Adding Sorenson's technique of applying an IP address blacklist (in addition to a domain name blacklist) to Treuhaft would improve Treuhaft in the same way it improves Sorenson's system. Sorenson receives a connection request including a domain name, compares the domain name against the blacklist, resolves the domain name into an IP address if the domain name is not on the blacklist, compares the resolved IP address against the IP address blacklist, and then allows the connection if the IP address is not on the blacklist. *See* Ex. 1009, ¶ [0031]-[0032], FIG. 6 (steps 608, 610, 612, 614, 616, 621, 622). Similarly, Treuhaft's name server receives a DNS query containing a domain name, resolves the domain name to its corresponding IP address, and applies the subscriber information to determine whether to return the resolved IP address or block that IP address and return a different one. Ex. 1008, ¶ [0064]-[0065]; FIG. 5B (steps 530-540).

227. In my opinion, because Treuhaft and Sorenson follow this same general sequence in processing a request, Sorenson's step of checking the IP address would be included as part of step 540 of Treuhaft's method 500 or added after step 540 as another step. *See id.*, ¶ [0065], FIG. 5B. In step 540, Treuhaft resolves the domain name in the DNS query to its corresponding IP address and determines whether to use that IP address or another IP address in the DNS response. *Id.*, ¶ [0065]. Having the resolved IP address, the DNS name server 120 of Treuhaft would simply check that IP address against Sorenson's added blacklist at that time as part of the decision whether to return it or return another address to the host device 105. Accordingly, in my opinion, Sorenson's known technique of checking the resolved IP address would be incorporated into Treuhaft's similar system to improve Treuhaft in the same way it benefits Sorenson.

228. In my opinion, Treuhaft stands ready for improvement by adding Sorenson's technique, and the combination would be made through routine skill in the art with predictable results. For example, Treuhaft's method 500 stands ready for improvement by adding Sorenson's IP address check as part of step 540, or another step following step 540, as discussed above. Since Treuhaft has already resolved the domain name into its IP address in step 540, the modified system would simply check the resolved IP address against the blacklist at that time. Thus, a POSA would incorporate Sorenson's technique without otherwise significantly modifying or redesigning Treuhaft's method.

229. Moreover, Treuhaft's DNS name server 120 already has a memory 200 storing subscriber information 280. *See* Ex. 1008, ¶¶ [0033]-[0034], FIG. 2. The subscriber information 280 likewise would be augmented to include Sorenson's blacklist information for each user or subscriber without otherwise significantly changing the structure or operation of the DNS name server 120. Thus, in my opinion, the combination would yield predictable results and would be made with a reasonable expectation of success. Accordingly, a POSA would have found it obvious to combine Treuhaft and Sorenson as proposed.

230. As explained above with respect to elements [1.3] and [1.4], in my opinion, Treuhaft discloses, or at least suggests, that its subscriber information 280 constitutes a preset list of authorized and/or unauthorized IP addresses. *Supra* Sections X.B.1.d, e. Thus, even though Sorenson discloses applying an IP address blacklist to discard the request if the resolved IP address is on the list--rather than "does not belong" to the list, as claimed, Treuhaft discloses or renders obvious applying a whitelist of authorized IP addresses, a blacklist of unauthorized IP addresses, or both in controlling the host device's access to the Internet. And, in the case of a whitelist of authorized IP addresses, Treuhaft discards the DNS query if the resolved IP address is not on the whitelist in the subscriber information 280. Thus, in my opinion, Treuhaft in combination with Sorenson renders obvious "discarding the service request packet if the [resolved] service server IP address does not belong to a preset service server IP address … in a preset list," as recited in elements [1.3], [6.3], and [11.3]." And for the same reasons, the combination of Treuhaft and Sorenson renders obvious that the preset list lists "service server IP addresses under an access authority of the terminal device" as recited in elements [1.4], [6.4], and [11.4], rather than not under an access authority of the terminal device. *Id.*

231. As discussed above, Treuhaft discloses or at least suggest implementing the subscriber information 280 as a "list" of authorized/unauthorized IP addresses. *Supra* Sections X.B.1.d, e. To the extent it is argued that Treuhaft alone does not expressly disclose or suggest a list, however, it is my opinion the combination with Sorenson's disclosure of an IP address blacklist further renders obvious using a "list" to implement Treuhaft authorized/unauthorized IP addresses. Accordingly, in my opinion, Treuhaft in view of Sorenson renders obvious elements [1.3], [1.4], [6.3], [6.4], [11.3], and [11.4], and so the combination of Treuhaft in view of Sorenson further renders obvious claims 1, 4-6, and 9-11 of the '040 Patent.

D. Grounds 5 and 6: Treuhaft/Sorenson in View of Bellinson Renders Obvious Claims 1, 4-6, and 9-11 of the '040 Patent Under § 103

232. Potentially, it may be argued that Treuhaft, or Treuhaft and Sorenson only disclose or suggest applying an IP address blacklist of unauthorized IP addresses, rather than a whitelist of authorized IP addresses, in determining whether to discard the DNS query. Thus, it may be argued, Treuhaft, or Treuhaft and Sorenson disclose or suggest discarding the service request packet if the resolved IP address is on the list, rather than "does not belong to a preset service server IP address … in a preset list," as recited in elements [1.3], [6.3], and [11.3]. For similar reasons, it may be argued Treuhaft's list contains IP addresses not under the authority of the host device 105, rather than "service server IP addresses under an access authority of the terminal device" as recited in elements [1.4], [6.4], and [11.4]. In opinion, such an argument would be incorrect at least for the reasons discussed above. *Supra* Section X.B.1.e.

233. Even if Patent Owner were correct, however, it is my opinion that Bellinson further discloses or suggests applying an allow-block list (i.e., a whitelist, blacklist, or combination of both) in determining whether to discard the request or establish a connection, and thus discloses these elements. As explained below, in my opinion, Treuhaft in view of Bellinson, and Treuhaft and Sorenson in view of Bellinson, further render obvious claims 1, 4-6, and 9-11 of the '040 Patent.

1. Bellinson Discloses Discarding the Request "if the [Resolved] Service Server IP Address Does Not Belong … in a Preset List" (Elements [1.3], [6.3], and [11.3])

234. In my opinion, Bellinson discloses elements [1.3], [6.3], and [11.3] by blocking a connection request if a requested site address is not found on an allow-block list (claimed preset

list). Specifically, Bellinson discloses "a system and method for controlling whether a user may access certain Internet sites" by applying "an allow-block list" to "determine[] whether the URL is referenced on the allow-block list and, if so, allow[] or disallow[] access to the site referenced by the URL accordingly." Ex. 1010, Abstract. In Bellinson, "[t]he allow-block list is a listing of specific site identifiers that the user is expressly authorized to view or prohibited from viewing." *Id.*, ¶ [0020]; *see also id.*, ¶¶ [0009] ("The allow-block list is a file containing a listing of specific URLs that the user is expressly authorized to view or expressly prohibited from viewing.").

235. As shown in Figure 3 (reproduced below), in step 244 Bellinson's system receives an access request containing "a specified site identifier that references an Internet site. Examples of such site identifiers include designators such as www.microsoft.com but could also include an Internet Protocol (IP) address." *Id.*, ¶ [0049], FIG. 3. In step 246 of Figure 3, Bellinson "determines whether the site identifier is on the allow-block list at step 246. *Id.* "If the site identifier is referenced on the allow-block list," in step 248 Bellinson "determine[s] whether the site identifier is designated as blocked on the allow-block list." *Id.*, ¶ [0050]. "If the site identifier is [designated as] blocked," Bellinson blocks the connection. *Id.* Otherwise, Bellinson allows the connection. *Id.*, ¶ [0051].

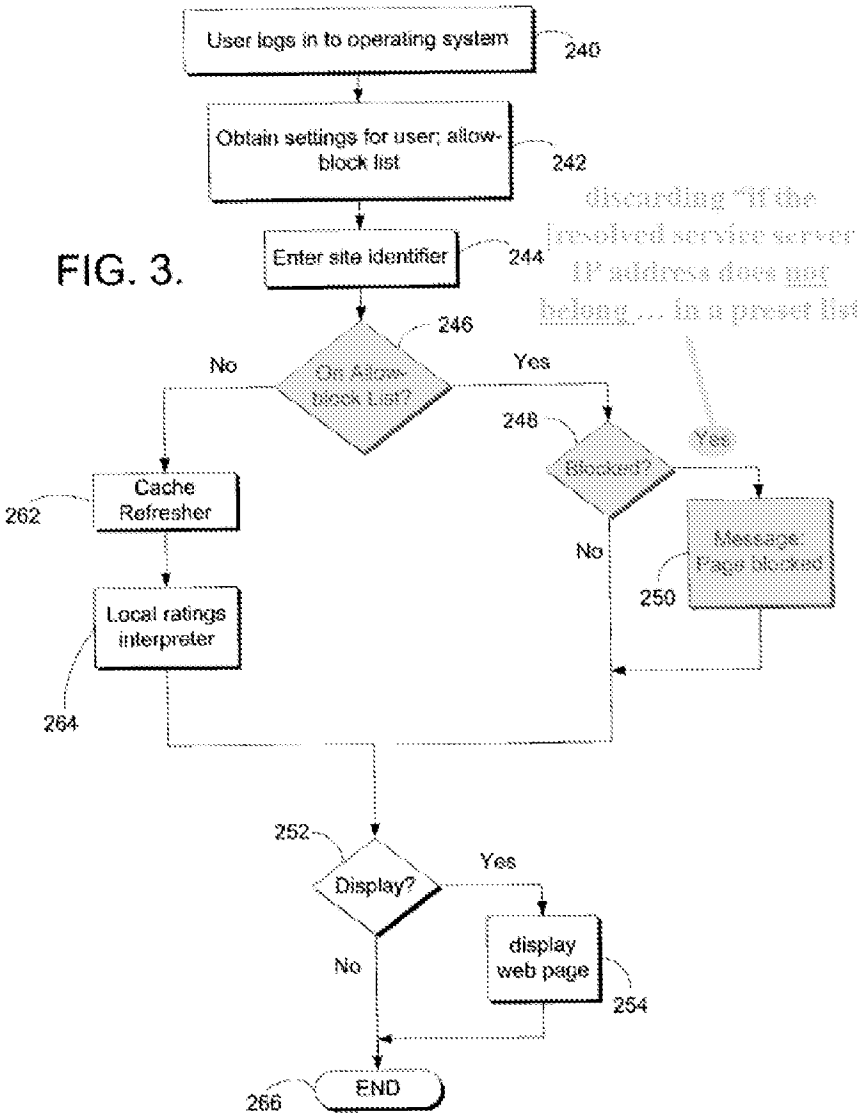

FIG. 3.

236. As highlighted in Figure 3, Bellinson discloses discarding the request "if the [resolved service server

IP address does not belong … in a preset list," as claimed, by determining that the allow-block list does not designate the site identifier as allowed in steps 246 and 248. *Id.*, ¶ [0050]. The allow-block list contains a list of site identifiers designated as allowed and a list of site identifiers designated as blocked. *Id.*, ¶¶ [0020] ("[t]he allow-block list is a listing of specific site identifiers that the user is expressly authorized to view or prohibited from viewing."), [0009]; Thus, in determining to block access to a site, Bellinson determines that the site identifier is not found on the allow portion of the list, because the block portion of the list designates the site identifier as blocked. In my opinion, by blocking the request if the address is not on the allow portion of the allow-block list, Bellinson discloses discarding the request "if the [resolved service server IP address does not belong … in a preset list," as recited in elements [1.3], [6.3], and [11.3].

2. Rationale to Combine Bellinson with Treuhaft and/or Treuhaft/Sorenson

237. In my opinion, a POSA would have found it obvious and been motivated to implement Bellinson's allow-block list in Treuhaft for several reasons. For example, the references provide teaching, suggestion, and/or motivation for making this combination. As explained above, Treuhaft's DNS name server 120 applies the subscriber information 280 for each user or subscriber to, among other things, block access to sites "categorized as an adult web site, a potential phishing or pharming site, and a website whose content has been deemed inappropriate by the user or containing material illegal in the country of the user." Ex. 1008, ¶ [0027]; *see also id.* ¶¶ [0006], [0028]. But Treuhaft does not block all connections—only those seeking access to unauthorized sites. Thus, in my opinion, Treuhaft contemplates both authorized and unauthorized sites, and Bellinson's allow-block list mechanism makes it possible to specifically designate both authorized and unauthorized sites. Accordingly, a POSA would have sought to incorporate Bellinson's allow-block list in Treuhaft as a mechanism to allow and block access to sites respectively deemed authorized and unauthorized.

238. Additionally, whitelists—like the allow portion of Bellinson's list, and blacklists—like the block portion of Bellinson's list, were known and used interchangeably and/or together in the same system to control access to sites. *See, e.g.*, Ex. 1011, 3[7]. Depending on the particular implementation, the administrator may find a whitelist, a blacklist, or a combination of both

[7] For ease of reference, Requester cites the PDF page number of Exhibit 1002.

appropriate for a given situation. *See, e.g.*, Ex. 1011, 3. For example, in some cases, the administrator may desire a strict approach to access control in which only connections to certain authorized sites are permitted, and thus may choose a whitelist. This approach may be useful, for example, in systems with child, student, and/or employee users that should only access certain specific websites or other resources. In other cases, the administrator may want to give users more leeway to access a variety of sites while blocking access to certain known unauthorized sites, and thus choose to use a blacklist. This approach may be useful, for example, if the administrator is more concerned with preventing malicious attacks on the user's computer rather than with preventing the user from accessing certain types of content.

239. The extent of an administrator's knowledge may also factor into the decision of whether to use a whitelist and/or a blacklist. *See* Ex. 1011, 3. For example, if the administrator has complete knowledge of the sites a user must or should access in the particular environment, a whitelist may be appropriate. But in cases where the administrator does not have such knowledge, a blacklist may be more appropriate. Accordingly, in my opinion, incorporating Bellinson's allow-block list into Treuhaft gives administrators more options for controlling users' or subscribers' access to the Internet, and thus would have sought to make the combination. Thus, a POSA would have had motivation to make the combination.

240. The combination of Bellinson with Treuhaft merely involves using a known technique (Bellinson's allow-block list) to improve a similar device (Treuhaft's system) in the same way and/or applying a known technique (Bellinson's allow-block list) to a known system (Treuhaft) ready for improvement to yield predictable results. Treuhaft and Bellinson describe systems similar in a number of ways. As discussed, Treuhaft allows an administrator to use subscriber information for users or subscribers of the system to block access to sites "categorized as an adult web site, a potential phishing or pharming site, and a website whose content has been deemed inappropriate by the user or containing material illegal in the country of the user" Ex. 1008, ¶ [0027]; *see also id.* ¶¶ [0006], [0028]. Similarly, Bellinson's technique using the allow-block list enables "parents to effectively control a child's web site access." Ex. 1010, ¶ [0008]; *see also id.*, ¶¶ [0003]-[0007], [0009], [0039]-[0041], [0048], [0057], [0058]. Thus, in my opinion, Treuhaft and Bellinson are analogous references with the same purpose of allowing administrators to prevent users from accessing inappropriate sites, services, or other online resources.

241. Treuhaft and Bellinson also have similar structure and operation. As discussed previously, Treuhaft's administrator inputs subscriber information 280 for a particular user or subscriber into the DNS name server 120, which the DNS name server 120 later applies to control that user's or subscriber's access to the Internet. *See, e.g.*, Ex. 1008, ¶¶ [0023], [0028], [0029], [0034], [0036], [0039], [0064], [0065]. Likewise, in Bellinson, "an administrator or parent would supply the content settings service with settings for a specified user," and those "settings could include the user's age group, age group map and an allow-block list." Ex. 1010, ¶ [0057]. Later, when that user sends a connection request including a site identifier, the Bellinson system retrieves the user's settings, including the allow-block list, and applies the allow-block list in determining whether to allow or block the connection to that site. *See, e.g., id.*, ¶¶ [0039], [0049]-[0051]. Accordingly, in my opinion, Treuhaft and Bellinson describe structurally- and functionally- similar systems designed to achieve a similar purpose.

242. Adding Bellinson's allow-block list to Treuhaft would improve Treuhaft in the same way it improves Bellinson's system. When receiving a connection request including a site identifier, Bellinson retrieves the user's previously-stored settings including the allow block list, Ex. 1010, ¶ [0039], FIG. 3 (step 242), and compares the site identifier to the allow-block list in determining whether to allow or block the connection to that site, *id.*, ¶¶ [0039], [0049]-[0051]. Similarly, when Treuhaft's DNS name server 120 receives a DNS query containing a domain name, it retrieves the previously stored subscriber information for the user or subscriber and applies the subscriber information in determining how to respond to the DNS query. Ex. 1008, ¶ [0064]-[0065]; FIG. 5B (steps 530-540).

243. In my opinion, because Treuhaft and Bellinson follow this same general sequence in processing a request, Bellinson's allow-blocklist would be incorporated into Treuhaft's process to improve Treuhaft in substantially the same way. For example, in the combined system, Bellinson's step of retrieving the user's settings including the allow-block list would be performed as part of Treuhaft's step 535 when the DNS name server 120 retrieves the user's subscriber information. *Compare* Ex. 1010, ¶ [0039], FIG. 3 (step 242) *with* Ex. 1008, ¶ [0064], FIG. 5B (step 535). Additionally, Bellinson's steps of applying the allow-block list would be performed as part of Treuhaft's corresponding step of applying the subscriber information to process the DNS query. *Compare* Ex. 1010, ¶¶ [0049]-[0051] , FIG. 3 (steps 246, 248) *with* Ex. 1008, ¶ [0065], [0066), FIG. 5B (steps 540, 545). Because the DNS name server 120 of Treuhaft has already resolved the

IP address at this point, it would simply check that IP address against Bellinson's allow-block list as part of the decision whether to return it or return another address to the host device 105. Accordingly, in my opinion, Bellinson's known technique of applying an allow-block list would be incorporated into Treuhaft's similar system to improve Treuhaft in the same way it benefits Sorenson.

244. Treuhaft stands ready for improvement by adding Bellinson's technique, and the combination would be made through routine skill in the art with predictable results. For example, Treuhaft's method 500 stands ready for improvement by adding Bellinson's allow-block list process, as discussed above. And since Treuhaft has already resolved the domain name into its IP address in step 540 or 545, the modified system would simply check the resolved IP address against the allow-block at that time. Thus, a POSA would incorporate Bellinson's technique without otherwise significantly modifying or redesigning Treuhaft's method.

245. Moreover, Treuhaft's DNS name server 120 already has a memory 200 storing subscriber information 280. *See* Ex. 1008, ¶¶ [0033]-[0034], FIG. 2. The subscriber information 280 likewise would be augmented to include Bellinson's allow-block list information for each user or subscriber without otherwise significantly changing the structure or operation of the DNS name server 120. Thus, in my opinion, the combination would yield predictable results and would be made with a reasonable expectation of success. *Id.* Accordingly, a POSA would have found it obvious to combine Treuhaft and Sorenson as proposed.

246. It is noted that the combination of Treuhaft and Bellinson additionally renders obvious the recitation in elements [1.4], [6.4], and [11.4] that the preset list contains "service server IP addresses under an access authority of the terminal device." Specifically, in my opinion, the allow portion of Bellinson's allow-block list, incorporated into Treuhaft, is a list of site identifiers under an access authority of Treuhaft's host device 105 (claimed terminal device).

247. As discussed above, Treuhaft discloses or at least suggests implementing the subscriber information 280 as a "list" of authorized/unauthorized IP addresses. *Supra* Sections X.B.1.d, e. But to the extent it is argued that Treuhaft alone does not expressly disclose or suggest a list, in my opinion, the combination with Bellinson's disclosure of an allow-block list further renders obvious using a "list" to implement Treuhaft authorized/unauthorized IP addresses.

248. Accordingly, for the reasons above, it is my opinion that the combination of Treuhaft in view of Bellinson and/or the combination of Treuhaft/Sorenson in view of Bellinson, render obvious elements [1.3], [1.4], [6.3], [6.4], [11.3], and [11.4]. Thus, these combinations of render obvious claims 1, 4-6, and 9-11 of the '040 Patent.

E. Secondary Considerations

249. This Declaration demonstrates that the Challenged Claims of the '040 Patent are unpatentable as anticipated and obvious in view of the prior art references. I have been told that the '040 Patent applicant did not identify any evidence of secondary considerations during prosecution. Further, I do not believe such secondary considerations would overcome the clear teachings in the references discussed above.

XI. CONCLUSION

250. For these reasons, it is my opinion that challenged claims 1, 4-6, and 9-11 of the '040 Patent are unpatentable.

I declare that all statements made herein of my knowledge are true, and that all statements made on information and belief are believed to be true, and that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: 11/08/2021 By: 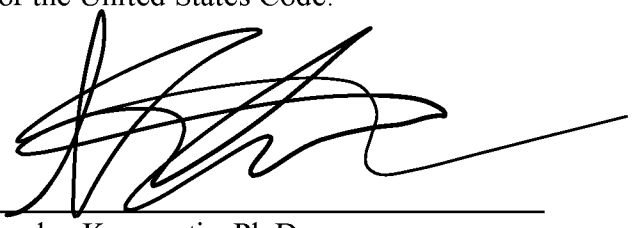

Angelos Keromytis, Ph.D.