

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

HOME DEPOT USA, INC.
Petitioner

v.


RAVENWHITE SECURITY, INC.
Patent Owner

Inter Partes Review No.: IPR2024-01316

**DECLARATION OF DR. CRAIG WILLS
REGARDING CLAIMS 1-10
OF U.S. PATENT NO. 10,594,823**

I hereby declare under penalty of perjury under the laws of the United States of America that the following is true and correct, and that all statements made of my own knowledge are true and that all statements made on information and belief are believed to be true. I understand that willful false statements are punishable by fine or imprisonment or both under Section 1001 of Title 18 of the United States Code.

Date: September 16, 2024



Dr. Craig Wills

TABLE OF CONTENTS

	Page
I. BACKGROUND AND QUALIFICATIONS	1
A. Qualifications	1
B. Technology Background.....	5
C. Materials Considered.....	15
II. LEGAL STANDARDS.....	15
A. Legal Standards for Prior Art	15
B. Legal Standards for Anticipation.....	16
C. Legal Standards for Obviousness.....	17
D. Legal Standards for Determining An Effective Filing Date	22
III. The '823 Patent.....	23
A. Overview of the '823 Patent.....	23
B. Overview of the '823 Patent's File History.....	28
C. Person of Ordinary Skill in the Art.....	30
D. Claim Construction Under 37 C.F.R. §§42.104(b)(3).....	30
E. Patent Owner's Infringement Contentions.....	32
IV. Ground 1: Hinton renders obvious claims 1-10	32
A. Overview of Hinton.....	32
B. Motivation to Combine	34
C. Claim 1	34
1. [1.pre] - A system, comprising:.....	34
2. [1.a] - one or more processors configured to	35
3. [1.b.i] – receive a network resource request from a client device wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session	36

4.	[1.b.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session.....	38
5.	[1.b.iii] - wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device	41
6.	[1.b.iv] - wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session	42
7.	[1.b.v] wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area.	45
8.	[1.c] - based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determine information that was encoded and stored in the client device;	46
9.	[1.d.i] - perform a first identification of at least one of the client device and a user of the client device using the first cookie of the first type,	48
10.	[1.d.ii] wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device	50
11.	[1.e] - perform a second identification of at least one of the client device and the user of the client device using the second cookie of the second type.....	51

12.	[1.f] - perform a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie	54
13.	[1.g] - a memory coupled to the one or more processors and configured to provide the one or more processors with instructions.....	55
D.	Claim 2 - The system of claim 1 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.....	56
E.	Claim 3 - The system of claim 1 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.....	57
F.	Claim 4 - The system of claim 1 wherein the determination comprises detection of pharming.	57
G.	Claim 5 - The system of claim 1 wherein in response to the performed determination, the one or more processors are configured to cause a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.	58
H.	Claim 6	59
1.	[6.pre] A method	60
2.	[6.a.i] receiving a network resource request from a client device, wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session,	60
3.	[6.a.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session,	60

4. [6.a.iii] wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device,60
 5. [6.a.iv] wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session,61
 6. [6.a.v] and wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area61
 7. [6.b] based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determining information that was encoded and stored in the client device; and61
 8. [6.c.i] performing a first identification of at least one of the client device and a user of the client device using the first cookie of the first type61
 9. [6.c.ii] wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device; and62
 10. [6.d] performing a second identification of at least one of the client device and the user of the client device using the second cookie of the second type; and.....62
 11. [6.e] performing a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request62
- I. Claim 7 - The method of claim 6 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.62

J.	Claim 8 - The method of claim 6 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.....	63
K.	Claim 9 - The method of claim 6 wherein the determination comprises detection of pharming.	63
L.	Claim 10 - The method of claim 6 wherein in response to the performed determination, further comprising causing a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.	63
V.	Ground 2: Varghese renders obvious claims 1, 3-6, And 8-10.....	64
A.	Overview of Varghese	64
B.	Motivation to Combine	68
C.	Claim 1	68
1.	[1.pre] - A system, comprising:.....	68
2.	[1.a] - one or more processors configured to	70
3.	[1.b.i] – receive a network resource request from a client device	70
4.	[1.b.ii] - wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session, wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session	71
5.	[1.b.iii] - wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device	74
6.	[1.b.iv] - wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session	75

7.	[1.b.v] - wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area.	77
8.	[1.c] - based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determine information that was encoded and stored in the client device;	79
9.	[1.d.i] and [1.d.ii] - perform a first identification of at least one of the client device and a user of the client device using the first cookie of the first type, wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device	80
10.	[1.e] - perform a second identification of at least one of the client device and the user of the client device using the second cookie of the second type.....	84
11.	[1.f] - perform a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie	86
12.	[1.g] - a memory coupled to the one or more processors and configured to provide the one or more processors with instructions.....	89
D.	Claim 3 - The system of claim 1 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.....	90
E.	Claim 4 - The system of claim 1 wherein the determination comprises detection of pharming.	91

F.	Claim 5 - The system of claim 1 wherein in response to the performed determination, the one or more processors are configured to cause a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type	92
G.	Claim 6	92
1.	[6.pre] A method, comprising	92
2.	[6.a.i] receiving a network resource request from a client device, wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session,	93
3.	[6.a.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session,	93
4.	[6.a.iii] wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device,	93
5.	[6.a.iv] wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session, and	93
6.	[6.a.v] wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area;	94
7.	[6.b] based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determining information that was encoded and stored in the client device; and	94

8.	[6.c.i] and [6.c.ii] performing a first identification of at least one of the client device and a user of the client device using the first cookie of the first type, wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device	94
9.	[6.d] performing a second identification of at least one of the client device and the user of the client device using the second cookie of the second type.....	94
10.	[6.e] performing a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie.	95
H.	Claim 8 - The method of claim 6 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.....	95
I.	Claim 9 - The method of claim 6 wherein the determination comprises detection of pharming.	95
J.	Claim 10 - The method of claim 6 wherein in response to the performed determination, further comprising causing a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.	95
VI.	Ground 3: Varghese In view of Hinton renders obvious claims 2 and 7.....	96
A.	Overview of Varghese	96
B.	Overview of Hinton.....	97
C.	Motivation to Combine	97
D.	Claim 2 - The system of claim 1 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.	98

E. Claim 7 - The method of claim 6 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.....99

VII. Secondary Considerations100

VIII. Availability for Cross-Examination.....100

IX. Right To Supplement100

X. Conclusion101

INDEX OF EXHIBITS

Exhibit No.	Description
1001	U.S. Patent No. 10,594,823 (the “ 823 patent ”).
1002	RESERVED
1003	File history of U.S. Patent No. 10,594,823.
1004	U.S. Patent No. 7,908,645 (“ Varghese ”).
1005	U.S. Patent Publication No. 2003/0115267 (“ Hinton ”)
1006	U.S. District Courts – Combined Civil and Criminal Federal Court Management Statistics (June 30, 2024).
1007	<i>Ravenwhite Licensing LLC v. The Home Depot, Inc.</i> , 2:24-cv-00688, Dkt. 1 (EDTX Aug. 21, 2024) (Complaint).
1008	<i>Ravenwhite Licensing LLC v. The Home Depot, Inc.</i> , 2:23-cv-00423, Infringement Contentions Cover Pleading (EDTX Dec. 4, 2023).
1009	<i>Ravenwhite Licensing LLC v. The Home Depot, Inc.</i> , 2:23-cv-00423, ‘823 Infringement Contentions (EDTX Dec. 4, 2023).
1010	Provisional Application No. 60/732,025 (“ Provisional ”)
1011	<i>Ravenwhite Licensing LLC v. The Home Depot, Inc.</i> , 2:23-cv-00423, Plaintiff’s P.R. 4-1 Disclosures (EDTX Aug. 8, 2024).
1012	Balachander Krishnamurthy and Craig E. Wills, <i>Generating a privacy footprint on the Internet</i> , In Proceedings of the ACM SIGCOMM Internet Measurement Conference, pages 65-70, Rio de Janeiro, Brazil (Oct. 2006).
1013	Jon Purdy, <i>Session Management for Clustered Applications</i> (Feb. 2005) (available at https://www.oracle.com/technical-resources/articles/enterprise-architecture/session-management.html).
1014	“Local Shared Objects—‘Flash Cookies,’” Electronic Privacy Information Center (EPIC) (July 21, 2005) (available at https://archive.epic.org/privacy/cookies/flash.html).
1015	“The Pharming Guide (part 2)” (Dec. 14, 2004) (available at: http://www.technicalinfo.net/papers/Pharming2.html).
1016	“Misfortune Cookies: Adjusting Internet Explorer to Block Tracking Web Cookies,” Gibson Research Corporation (last modified Aug. 13, 2005) (available at https://www.grc.com/cookies.htm).
1017	Ileene Chernoff, “Cookie crumbs, an introduction to cookies,” SANS Institute (2005) (available at https://www.giac.org/paper/gsec/226/cookie-crumbs-introduction-cookies/100727).

Exhibit No.	Description
1018	Michael Nelte and Elton Saul, "Cookies: Weaving the Web into a State," Crossroads, The ACM Magazine for Students, Vol. 7, Issue 1, pp. 10-13, (Sept. 1, 2000) (available at https://dl.acm.org/doi/10.1145/351092.351097).
1019	Edward W. Felten and Michael A. Schneider, "Timing Attacks on Web Privacy" (Nov. 25, 2002) (available at https://web.archive.org/web/20021125051243/http://www.cs.princeton.edu/sip/pub/webtiming.pdf).
1020	SecuriTeam.com, "Timing Attacks of Web Privacy (Paper and Specific Issue)" (Feb. 20, 2002) (available at https://web.archive.org/web/20021020062537/http://www.securiteam.com/securityreviews/5GP020A6LG.html).
1021	Martin Pool, "meantime: non-consensual http user tracking using caches" (last revised March 29, 2000) (available at https://sourcefrog.net/projects/meantime).
1022	United Virtualities, "United Virtualities Develops ID Backup to Cookies" (March 31, 2005) (available at https://web.archive.org/web/20050408075600/http://www.unitedvirtualities.com/UV-Pressrelease03-31-05.htm).
1023	Antone Gonsalves, "Company Bypasses Cookie-Deleting Consumers," Information Week (March 31, 2005) (available at https://www.informationweek.com/it-leadership/company-bypasses-cookie-deleting-consumers).
1024	Dr. Craig Wills, CS3013 Course Notes Week 4 (2004) (available at https://web.cs.wpi.edu/~cs3013/c04/week4-memmgmt.pdf).
1025	Prof. Howard Hamilton, CS 330 Course Notes (2003) (available at https://www2.cs.uregina.ca/~hamilton/courses/330/notes/memory/MemoryHierarchy.html).
1026	John Schwartz, "Giving Web a Memory Cost Its Users Privacy," New York Times (Sept. 4, 2001) (available at https://www.nytimes.com/2001/09/04/business/giving-web-a-memory-cost-its-users-privacy.html).
1027	Thomas Chung, "HOWTO: Installing Flash Plugin in Firefox Way," FedoraNEWS.ORG (Sept. 16, 2004) (available at https://fedoranews.org/tchung/firefox-flash/).
1028	Adrian Ludwig, "Macromedia® Flash® Platform Security and Macromedia Enterprise Solutions" (Sept. 2005) (available at

Exhibit No.	Description
	https://www.adobe.com/platform/whitepapers/flashplatform_security_enterprise.pdf).

CHART OF CLAIMS

[1.pre] A system, comprising:
[1.a] one or more processors configured to:
[1.b.i] receive a network resource request from a client device, wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session,
[1.b.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session,
[1.b.iii] wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device,
[1.b.iv] wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session, and
[1.b.v] wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area;
[1.c] based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determine information that was encoded and stored in the client device;
[1.d.i] perform a first identification of at least one of the client device and a user of the client device using the first cookie of the first type,
[1.d.ii] wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device;
[1.e] perform a second identification of at least one of the client device and the user of the client device using the second cookie of the second type; and
[1.f] perform a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie; and
[1.g] a memory coupled to the one or more processors and configured to provide the one or more processors with instructions.
[2] The system of claim 1 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to

be stored to the client device during the first and second previous network sessions.
[3] The system of claim 1 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.
[5] The system of claim 1 wherein in response to the performed determination, the one or more processors are configured to cause a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.
[6.pre] A method, comprising:
[6.a.i] receiving a network resource request from a client device, wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session,
[6.a.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session,
[6.a.iii] wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device,
[6.a.iv] wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session, and
[6.a.v] wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area;
[6.b] based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determining information that was encoded and stored in the client device; and
[6.c.i] performing a first identification of at least one of the client device and a user of the client device using the first cookie of the first type
[6.c.ii] wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device; and
[6.d] performing a second identification of at least one of the client device and the user of the client device using the second cookie of the second type; and
[6.e] performing a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second

cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie.

[7] The method of claim 6 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.

[8] The method of claim 6 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.

[10] The method of claim 6 wherein in response to the performed determination, further comprising causing a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.

I, Craig Wills, Ph.D., declare as follows:

1. My name is Craig Wills, and I have been retained by counsel for Home Depot USA, Inc. (“Petitioner”) to analyze U.S. Patent No. 10,594,823 (“’823 patent”) and to provide my opinions regarding the patentability of claims 1-10 of the ’823 patent.

2. I am being compensated at my normal consulting rate of \$400 per hour for my time. My compensation is not contingent on the outcome of this proceeding or of any proceedings relating to the ’823 patent.

I. BACKGROUND AND QUALIFICATIONS

A. Qualifications

3. I am a Professor of Computer Science at Worcester Polytechnic Institute (“WPI”). One of my areas of expertise is Web privacy where I have a number of peer-reviewed publications, have given many invited presentations and have had my work appear in venues such as the Wall Street Journal, the New York Times and National Public Radio.

4. I received my B.S. in computer science from University of Nebraska (1982) as well as my M.S. (1984) and Ph.D. (1988) in computer science from Purdue University. Before my appointment at WPI, I worked at AT&T Bell Laboratories where my work focused on the design and development of an automation tool for network management applications. Since starting at WPI, I

have had visiting positions with Cisco Systems, Inc., where I worked with the Network Management Technology Group, at the School of Mathematical and Computing Sciences at Victoria University of Wellington, and at the University of Bologna.

5. My teaching and research span many areas of computer science, including Internet application performance, distributed and mobile computing, networking, and Web privacy. My work has been published in refereed journals and presented at top research conferences. My work also includes research on social networks, mobile social networks and available applications such as network games. My work has been cited in popular press venues such as the *New York Times*, *Wall Street Journal*, *USA Today*, *Los Angeles Times*, *San Jose Mercury News*, *Atlanta Journal-Constitution*, *InformationWeek*, and *National Public Radio Science Friday*.

6. I am a named inventor on eight patents, including U.S. Patent No. 9,172,706 entitled “Tailored protection of personally identifiable information,” U.S. Patent No. 8,601,591 entitled “Method and apparatus for providing web privacy,” and U.S. Patent No. 7,296,089 entitled “Method for improving web performance by adapting servers based on client cluster characterization.”

7. I am a member of the Association for Computing Machinery (ACM) and the IEEE Computer Society. At ACM, I was a founding associate editor of

ACM Transactions on Internet Technology, where I served from July 2000-July 2009.

8. My own research considered the use of Web cookies before and around the time of the '823 patent. Such publications include:

- Balachander Krishnamurthy and Craig E. Wills. Proxy cache coherency and replacement—towards a more complete picture. In Proceedings of the 19th IEEE International Conference on Distributed Computing Systems, pages 332–339, Austin, TX, June 1999.
- Balachander Krishnamurthy and Craig Wills. Cat and mouse: Content delivery tradeoffs in web access. In Proceedings of the International World Wide Web Conference, pages 337–346, Edinburgh, Scotland, May 2006.
- Balachander Krishnamurthy and Craig E. Wills. Generating a privacy footprint on the Internet. In Proceedings of the ACM SIGCOMM Internet Measurement Conference, pages 65–70, Rio de Janeiro, Brazil, October 2006. EX1012.

9. My own subsequent work made direct use of different types of cookies:

First/third-party cookies:

- Balachander Krishnamurthy, Delfina Malandrino, and Craig E. Wills. Measuring privacy loss and the impact of privacy protection in web

browsing. In Proceedings of the Symposium on Usable Privacy and Security, pages 52–63, Pittsburgh, PA USA, July 2007.

- Craig E. Wills. Identifying and preventing conditions for web privacy leakage. In Proceedings of the *W3C Workshop on Web Tracking and User Privacy*, pages 1–5, Princeton, NJ USA, April 2011.

“CSS” history-based cookies:

- Craig E. Wills and Mihajlo Zeljkovic. A personalized approach to web privacy—awareness, attitudes and actions. *Information Management and Computer Security*, 19(1):53–73, 2011.

Flash cookies (as well as first/third-party cookies):

- Balachander Krishnamurthy, Konstantin Naryshkin, and Craig E. Wills. Privacy leakage vs. protection measures: The growing disconnect. In Proceedings of the Web 2.0 Security and Privacy Workshop, pages 1–10, Oakland, CA USA, May 2011.
- Craig E. Wills and Can Tatar. Understanding what they do with what they know. In Proceedings of the Workshop on Privacy in the Electronic Society, Raleigh, NC USA, October 2012.

10. My *curriculum vitae*, attached as Appendix 1, contains a more detailed description of my background.

B. Technology Background

11. The technical area of the '823 patent is the use of different types of cookies for a Web browser running on a client machine to store information in a browser storage area (browser cache), which has been sent to it by a (Web) server over the Internet. The stored cookie information in the storage area is sent by the client to the server on subsequent Web requests.

12. This original type of cookies can be described as “Cookies are small text files used to save information about an individual or their use of a web site. For instance, a cookie can be used to save your login name, your preferences for viewing content, or to track you as you browse the Internet.” EX1014 (“Local Shared Objects—‘Flash Cookies’” published July 21, 2005 from the Electronic Privacy Information Center (EPIC) available at <https://archive.epic.org/privacy/cookies/flash.html> and referenced as one of the Other Publications listed on the face of the '823 patent).

13. Identification of different types of cookies for Web interactions between a browser and a server was well-known at the time of the patent. As more cookie types have been identified we do observe that the original use of a “Web cookie” has been qualified with a variety of descriptors such as normal, traditional, conventional, standard and regular in the literature. Each of these qualifiers refer to

the same original cookie type. The identified cookie types at the time of the '823 patent include:

14. A 2001 article in the New York Times both describes the development of the original Web cookie as well as public reaction that led to the distinction of two cookie types based on the Internet domain from which they are served. “By 1996, the existence of cookies and third-party cookies was becoming a hot topic in the news media and in online forums; Mr. Montulli and Netscape altered the company’s browsers to distinguish **cookies coming directly from the site being viewed [first-party cookies]** from **third-party cookies** and to give consumers some control over them, allowing them to turn off all cookies or just the third-party variety. Microsoft, too, implemented some cookie control tools over time. But by default, browsers were set (and are still set) to accept such cookies automatically unless the user told the software not to -- which meant that a great majority of people ended up accepting cookies unknowingly from nearly every site they had visited.” EX1026 (“Giving Web a Memory Cost Its Users Privacy” by John Schwartz, The New York Times, September 4, 2001 available at <https://www.nytimes.com/2001/09/04/business/giving-web-a-memory-cost-its-users-privacy.html>). The Gibson Research Corporation has a Misfortune Cookies web page from August 2005 (EX1016) (available at <https://www.grc.com/cookies.htm>) that also distinguishes between first-party and

third-party cookies. Regarding first-party cookies it describes: “So called “first-party” cookies are those offered to your web browser by the same web site you are visiting. While these cookies—whether they are session-cookies or persistent-cookies—do allow your movement around the site to be followed by the site, they are not generally regarded as a privacy concern because they can not be seen or accessed by other sites.” In contrast, it describes third-party cookies as the “real problem” from a privacy standpoint: “Part of the amazing power of the web is that a page can instruct your web browser to fetch any other pieces of the web page from any other servers located anywhere on the Internet. This is a powerful capability. But it comes at some cost of privacy and security because you, the trusting web user, are unable to exercise any control over which other “third-party” Internet servers your own web browser will be connecting to and requesting data from.”

15. Cookies are typed not only based upon the party in which they are served from, but also based on the presence or absence of an expiration time. A Global Information Assurance Certification paper “Cookie crumbs, an introduction to cookies” (EX1017) (available at <https://www.giac.org/paper/gsec/226/cookie-crumbs-introduction-cookies/100727> with Google search showing a date of February 12, 2000) describes how this distinction results in two types of cookies being created—“A persistent or session specific cookie may be created. **Persistent**

cookies are stored on one’s hard drive, where **session cookies** are temporary and are deleted when your browser is shut down or logged off.” Whether a Web cookie is persistent or session specific is determined by the Expires attribute of the Set-Cookie response header. If this attribute is explicitly set then the cookie is persistent otherwise it expires at the end of the current browsing session (EX1018) (“Cookies: Weaving the Web into a State” by Michael Nelte and Elton Saul, Crossroads The ACM Student Magazine, 2000). The previously cited Misfortunate Cookies Web page also distinguishes between persistent and session cookies and notes how each are stored in the browser: “A ‘session cookie’ is a non-persistent cookie which a web browser agrees to accept and carry—but only for the current web surfing session. Unlike regular cookies which are persistent and can be retained and carried in a browser for years, session cookies are kept in memory and are not written to the system’s permanent storage. They are discarded when the browser or computer system is shut down.”

16. In a 2000 paper on “Timing Attacks on Web Privacy” (EX1019) (available at <https://web.archive.org/web/20021125051243/http://www.cs.princeton.edu/sip/pub/webtiming.pdf>), Felten and Schneider introduce a new type of cookie, which is a file, such as an image, being written into the browser cache without knowledge of the user: “Using the methods described above, it is possible to develop a new, more

intrusive, form of web cookies, which we call ‘**cache cookies**’. Servers can store cache cookies on clients who visit their pages, without the client’s knowledge or approval. Cache cookies are stored in the form of entries in the client’s web cache. By forcing a client to retrieve a specific URL (using an IMG tag, for example), the server can effectively write entries into the client’s cache, thus storing the cookie. To read the cookie, the server can use any of the measurement techniques described above to measure the retrieval time for the same URL. A hit indicates that the cookie is present, otherwise the cookie is not present.”

17. In 2002, Clover (EX1020) (available at <https://web.archive.org/web/20021020062537/http://www.securiteam.com/security/reviews/5GP020A6LG.html>) follows up the Felten and Schneider paper by identifying another cookie type noting that “CSS [Cascading Style Sheets] has a feature that can be abused to exactly the same ends. It is simpler, more accurate, and more easily abused than the timing attacks described in the above paper.” Clover indicates the specific issue that “The CSS :visited pseudo-class can be used to apply different on-screen styling to links leading to pages the user has already visited.” Clover goes on to describe how this feature can be used for “**CSS cookies**” based on the history of visited links maintained by a browser.

18. The previously-cited 2005 article from EPIC (EX1014) references a study on user practices with standard cookies where a majority reported having

deleted cookies from their computer. This behavior led companies to employ yet another type of cookie. Specifically, the article reports “With the advent of spyware and spyware removal programs, as well as media attention and the increase of online literacy, users now understand the purposes and risks of using cookies. Recently, users have become more vigilant in purging cookies from their computers. According to a Jupiter Research study, 58% of online users have deleted cookies from their computer and 39% of users do so on a monthly basis. This regular ‘cookie tossing’ is causing direct marketers to see more invasive methods to track individuals. One of those methods is to set a ‘Local Shared Object,’ also known as a ‘**Flash cookie**’ to track individuals. Simply put, the idea behind this tracking is to set two cookies on the user’s machine--a standard cookie that the consumer may erase, and a second Flash cookie that the user probably will keep, because the existence of Flash cookies is not well known.”

19. While not a separate type of cookie, a blog entry from 2000 describes a privacy attack that exploits HTTP cache-control headers such as If-Modified-Since and Etag as a means to assign unique identifiers to “allow servers to track individual users in a manner similar to cookies, but with less constraints.” (EX1021) (“meantime: non-consensual http user tracking using caches”, by Martin Pool, 2000 available at <https://sourcefrog.net/projects/meantime>).

20. Thus, at the time of the '823 patent, a number of cookie types were known. The original cookie type was further broken down into types based on whether a cookie is first-party/third-party or persistent/session. New cookie types, such as cache cookies, CSS cookies and Flash cookies were also being identified. The use of multiple cookie types was of interest in a range of domains including advertising, cross-domain user identification, and identity-theft prevention.

21. By 2005 the use of various cookie types had become a concern of both advertisers and privacy advocates. Into this mix the company United Virtualities announced in March, 2005 a browser-based “Persistent Identification Element” allowing erased cookies to be restored. A portion of the press release (EX1022) (available at <https://web.archive.org/web/20050408075600/http://www.unitedvirtualities.com/UV-Pressrelease03-31-05.htm>) describes:

NEW YORK (March 31, 2005) United Virtualities, the leading innovator of creative marketing and technology solutions for the digital marketplace, today announced it has developed a backup ID system for cookies set by web sites, ad networks and advertisers, but increasingly deleted by users. UV’s “Persistent Identification Element” (PIE) is tagged to the user’s browser, providing each with a unique ID just like traditional cookie coding. However, PIEs cannot be deleted by any commercially available anti-

spyware, mal-ware, or adware removal program. They will even function at the default security setting for Internet Explorer.

22. “All advertisers, websites and networks use cookies for targeted advertising, but cookies are under attack. According to current research they are being erased by 40% of users creating serious problems,” says Mookie Tenenbaum, founder of United Virtualities. “From simple frequency capping to the more sophisticated behavioral targeting, cookies are an essential part of any online ad campaign. PIE will give publishers and third-party providers a persistent backup to cookies effectively rendering them unassailable.”

23. There are two types of PIES:

- AccuCounter PIE, a cookie replacement that counts unique users accurately.
- Backup PIE: a PIE that not only counts unique users but also recognizes the visitor and restores any erased cookies.

24. “The erasing of cookies threatens many cookie dependent server-side applications from registration to targeting to traffic counting,” says Mr. Tenenbaum. “PIES are a cookie support product that ensures persistent identification of the users.”

25. Implementation of the PIE technology is instantaneous and requires the insertion of just a line of code. UV plans to sell the PIE technology to publishers and Networks who are worried about cookies being deleted.

26. There was much news coverage of this announcement with more details on how this technology works and its implications. One such article (EX1023) (“Company Bypasses Cookie-Deleting Consumers” by Antone Gonsalves, Information Week, March 31, 2005, available at <https://www.informationweek.com/it-leadership/company-bypasses-cookie-deleting-consumers>) describes:

United Virtualities’s PIE helps combat this consumer behavior by leveraging a feature in Flash MX called local shared objects. Flash MX is a Macromedia Inc. application for developing multimedia Web content, user interfaces and Web applications. The technology runs on a Flash Player that the company says is deployed on 98 percent of Internet-capable computers.

27. When a consumer goes to a PIE-enabled website, the visitor’s browser is tagged with a Flash object that contains a unique identification similar to the text found in a traditional cookie. In this way, PIE acts as a cookie backup, and can also restore the original cookie when the consumer revisits the site.

28. While consumers have learned to delete cookies, most are unaware of shared objects, and don’t know how to disable them.

29. Similarly, the 2005 EPIC article on Flash cookies (EX1014) addresses the use of the Persistent Identification Element (PIE):

United Virtualities (UV), an online marketing firm, has introduced a tracking platform that takes advantage of the relative obscurity of Flash cookies. In a press release this March, UV announced PIE, a backup ID system for cookies. Mookie Tenenbaum, founder of United Virtualities, explained the reasoning behind the product, “All advertisers, websites and networks use cookies for targeted advertising, but cookies are under attack. According to current research they are being erased by 40% of users creating serious problems.”

UV’s press release also claims that the PIE system can restore deleted web cookies. Although there is little official information on the implementation of the PIE system, it is not likely that the cookie is actually restored. Instead, it appears that the Flash cookie acts as a redundancy. That is, the PIE system uses Flash cookies as a backup. A site interested in tracking a user would set a normal cookie and a Flash cookie. If the user erased the normal cookie, the PIE-enabled site could use the redundant Flash cookie to track the user.

To justify this tracking mechanism, UV’s Tenenbaum said, “The user is not proficient enough in technology to know if the cookie is good or bad, or how it works.”

This practice is highly deceptive. By deleting cookies, consumers are clearly rejecting attempts to track them. Using an obscure technology to subvert these wishes is a practice that should be stopped. Cookies have many beneficial purposes and can make the end user's web experience better. Websites should be honest and up front about how they use cookies, and they should respect the decisions of those users who do not want to be tracked via cookies.

EX1014.

C. Materials Considered

30. In preparing this declaration, I have reviewed and/or considered at least the documents cited in the List of Exhibits, and the documents referenced in this declaration. I have also reviewed the prosecution history of the '823 patent.

II. LEGAL STANDARDS

A. Legal Standards for Prior Art

31. I understand that a patent or other publication must first qualify as prior art before it can be used to invalidate a patent claim.

32. I understand that a U.S. or foreign patent qualifies as prior art to an asserted patent if the date of issuance of the patent is prior to the invention of the asserted patent. I further understand that a printed publication, such as an article

published in a magazine or trade publication, qualifies as prior art to an asserted patent if the date of publication is prior to the invention of the asserted patent.

33. I understand that a U.S. or foreign patent also qualifies as prior art to an asserted patent if the date of issuance of the patent is more than one year before the filing date of the asserted patent. I further understand that a printed publication, such as an article published in a magazine or trade publication, constitutes prior art to an asserted patent if the publication occurs more than one year before the filing date of the asserted patent.

34. I understand that a U.S. patent qualifies as prior art to the asserted patent if the application for that patent was filed in the United States before the invention of the asserted patent.

B. Legal Standards for Anticipation

35. I understand that documents and materials that qualify as prior art can be used to invalidate a patent claim via anticipation or obviousness.

36. I understand that, once the claims of a patent have been properly construed, the second step in determining anticipation of a patent claim requires a comparison of the properly construed claim language to the prior art on a limitation-by-limitation basis.

37. I understand that a prior art reference “anticipates” an asserted claim, and thus renders the claim invalid, if all elements of the claim are disclosed in that prior art reference, either explicitly or inherently (*i.e.*, necessarily present).

38. I understand that anticipation in an *inter partes* review must be shown by a preponderance of the evidence.

C. Legal Standards for Obviousness

39. I understand that even if a patent is not anticipated, it is still invalid if the differences between the claimed subject matter and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person of ordinary skill in the pertinent art.

40. I understand that a person of ordinary skill in the art provides a reference point from which the prior art and claimed invention should be viewed. This reference point prevents one from using his or her own insight or hindsight in deciding whether a claim is obvious.

41. I also understand that an obviousness determination includes the consideration of various factors such as (1) the scope and content of the prior art, (2) the differences between the prior art and the asserted claims, (3) the level of ordinary skill in the pertinent art, and (4) the existence of secondary considerations such as commercial success, long-felt but unresolved needs, failure of others, etc.

42. I understand that an obviousness evaluation can be based on a combination of multiple prior art references. I understand that the prior art references themselves may provide a suggestion, motivation, or reason to combine, but other times the nexus linking two or more prior art references is simple common sense. I further understand that obviousness analysis recognizes that market demand, rather than scientific literature, often drives innovation, and that a motivation to combine references may be supplied by the direction of the marketplace.

43. I understand that if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.

44. I also understand that practical and common-sense considerations should guide a proper obviousness analysis, because familiar items may have obvious uses beyond their primary purposes. I further understand that a person of ordinary skill in the art looking to overcome a problem will often be able to fit together the teachings of multiple publications. I understand that obviousness analysis therefore takes into account the inferences and creative steps that a person of ordinary skill in the art would employ under the circumstances.

45. I understand that a particular combination may be proven obvious merely by showing that it was obvious to try the combination. For example, when there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp because the result is likely the product not of innovation but of ordinary skill and common sense.

46. The combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results. When a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, the patent claim is likely obvious.

47. It is further my understanding that a proper obviousness analysis focuses on what was known or obvious to a person of ordinary skill in the art, not just the patentee. Accordingly, I understand that any need or problem known in the field of endeavor at the time of invention and addressed by the patent can provide a reason for combining the elements in the manner claimed.

48. I understand that a proposed obviousness combination does not need to be the preferred, or most desirable, combination available, in order to render a

claim obvious. Instead, there must be something that suggests the desirability of the proposed combination, not necessarily something that suggests that the proposed combination is the most desirable combination available.

49. Similarly, I understand that for purposes of obviousness, when a motivating benefit comes at the expense of another benefit, that does not nullify the proposed obviousness combination or modification. Likewise, a proposed obviousness combination does not fail merely because it may be inferior to other systems used for the same use.

50. I also understand that a determination of obviousness based on teachings from multiple references does not require an actual, physical substitution of elements between the multiple references. Likewise, the test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference. Instead, the test is whether a claim is rendered obvious by the teachings of the prior art as a whole.

51. I understand that a claim can be obvious in light of a single reference, without the need to combine references, if the elements of the claim that are not found explicitly or inherently in the reference can be supplied by the common sense of one of skill in the art.

52. I understand that secondary indicia of non-obviousness may include:
(1) a long felt but unmet need in the prior art that was satisfied by the invention of

the patent; (2) commercial success of processes covered by the patent; (3) unexpected results achieved by the invention; (4) praise of the invention by others skilled in the art; (5) taking of licenses under the patent by others; (6) deliberate copying of the invention; (7) failure of others to find a solution to the long felt need; and (8) skepticism by experts.

53. I also understand that there must be a relationship between any such secondary considerations and the invention. I further understand that contemporaneous and independent invention by others is a secondary consideration supporting an obviousness determination.

54. In sum, my understanding is that prior art teachings are properly combined where a person of ordinary skill in the art having the understanding and knowledge reflected in the prior art and motivated by the general problem facing the inventor, would have been led to make the combination of elements recited in the claims. Under this analysis, the prior art references themselves, or any need or problem known in the field of endeavor at the time of the invention, can provide a reason for combining the elements of multiple prior art references in the claimed manner.

55. I understand that obviousness in an *inter partes* review must be shown by a preponderance of the evidence.

D. Legal Standards for Determining An Effective Filing Date

56. Petitioner’s counsel has informed me that “[i]t is elementary patent law that a patent application is entitled to the benefit of the filing date of an earlier filed application only if the disclosure of the earlier application provides support for the claims of the later application, as required by 35 U.S.C. §112.” *In re Chu*, 66 F.3d 292, 297 (Fed. Cir. 1995); 35 U.S.C. §111. “In other words, the specification of the *provisional* must ‘contain a written description of the invention and the manner and process of making and using it, in such full, clear, concise, and exact terms,’ 35 U.S.C. § 112 ¶ 1, to enable an ordinarily skilled artisan to practice the invention *claimed* in the *non-provisional* application.” *New Railhead Mfg., LLC v. Vermeer Mfg. Co.*, 298 F. 3d 1290, 1294 (Fed. Circ. 2002) (emphasis in original). When a priority claim involves a chain of priority documents, “each application in the chain leading back to the earlier application must comply with the written description requirement of 35 U.S.C. §112.” *Lockwood v. Am. Airlines, Inc.*, 107 F.3d 1565, 1572 (Fed. Cir. 1997). To satisfy the written description requirement of 35 U.S.C. §112, the disclosure of the earlier filed application must “convey with reasonable clarity to those skilled in the art that, as of the filing date sought, [the inventor] was in possession of *the invention*.” *Vas-Cath Inc. v. Mahurkar*, 935 F.2d 1555, 1563–64 (Fed. Cir. 1991) (emphasis in original). A prior application that merely renders the later-claimed invention obvious is not

sufficient to meet the written description requirement—it must describe the claimed invention with all its limitations. *Tronzo v. Biomet, Inc.*, 156 F.3d 1154, 1158 (Fed. Cir. 1998). Similarly, “[i]t is not sufficient for purposes of the written description requirement of §112 that the disclosure, when combined with the knowledge in the art, would lead one to speculate as to modifications that the inventor might have envisioned, but failed to disclose.” *Lockwood*, 107 F.3d at 1572.

III. THE '823 PATENT

A. Overview of the '823 Patent

57. The '823 patent issued from U.S. Patent Application No. 15/706,556, filed September 15, 2017. The '823 patent claims priority to a series of applications, the earliest of which is U.S. Patent Application Serial No. 11/590,083, filed October 31, 2006. I have been informed that, even though the '823 patent lists a priority claim to an earlier-filed provisional application, in the underlying litigation, Patent Owner only asserted a priority claim of October 31, 2006. EX1008, 5.

58. I have been informed that “[i]t is elementary patent law that a patent application is entitled to the benefit of the filing date of an earlier filed application only if the disclosure of the earlier application provides support for the claims of the later application, as required by 35 U.S.C. §112.” *In re Chu*, 66 F.3d 292, 297

(Fed. Cir. 1995). When a priority claim involves a chain of priority documents, “each application in the chain leading back to the earlier application must comply with the written description requirement of 35 U.S.C. §112.” *Lockwood v. Am. Airlines, Inc.*, 107 F.3d 1565, 1571 (Fed. Cir. 1997). To satisfy the written description requirement of 35 U.S.C. §112, the disclosure of the earlier filed application must “convey with reasonable clarity to those skilled in the art that, as of the filing date sought, [the inventor] was in possession of *the invention*.” *Vas-Cath Inc. v. Mahurkar*, 935 F.2d 1555, 1563-64 (Fed. Cir. 1991) (emphasis in original). “In other words, the specification of the *provisional* must ‘contain a written description of the invention and the manner and process of making and using it, in such full, clear, concise, and exact terms,’ 35 U.S.C. § 112 ¶ 1, to enable an ordinarily skilled artisan to practice the invention *claimed* in the *non-provisional* application.” *New Railhead Mfg., LLC v. Vermeer Mfg. Co.*, 298 F. 3d 1290, 1294 (Fed. Circ. 2002) (emphasis in original). A prior application that merely renders the later-claimed invention obvious is not sufficient to meet the written description requirement—it must describe the claimed invention with all its limitations. *Tronzo v. Biomet, Inc.*, 156 F.3d 1154, 1158 (Fed. Cir. 1998). Similarly, “[i]t is not sufficient for purposes of the written description requirement of § 112 that the disclosure, when combined with the knowledge in the art, would lead one to speculate as to modifications that the inventor might have envisioned, but failed to

disclose.” *Lockwood*, 107 F.3d at 1572. I have been informed that even when “each element may be individually described in the specification, ... the lack of adequate description of their combination” renders those claimed combinations invalid under Section 112. *Hyatt v. Dudas*, 492 F.3d 1365, 1377 (Fed. Cir. 2007) (upholding the patent examiner’s finding that “it is not enough that applicant show where each claimed element resides in the earliest filed application but [he] must also provide support for the linkage of the claimed elements creating the embodiment.”) (emphasis in original); *Flash-Control, LLC v. Intel Corp.*, No. 2020-2141, 2021 WL 2944592, at *4 (Fed. Cir. July 14, 2021) (“A patent owner cannot show written description support by picking and choosing claim elements from different embodiments that are never linked together in the specification.”); *Novozymes A/S v. DuPont Nutrition Biosciences APS*, 723 F.3d 1336, 1349 (Fed. Cir. 2013) (“The specification must present each claim as an ‘integrated whole.’”); *Ariad Pharms., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336, 1352 (Fed. Cir. 2010) (“[A] description that merely renders the invention obvious does not satisfy the [written description] requirement.”).

59. I have reviewed the Provisional (EX1010). In my opinion, the provisional does not provide written description support for the claims of the ’823 patent. For example, the Provisional does not include disclosure of elements 1.b.ii, 1.b.iii, 1.b.iv, 6.a.ii, or 6.a.iii. EX1010. These limitations were added by

amendment based on paragraph 30 of the specification (EX1003, 226, 232), but that paragraph or similar disclosure is not in the Provisional. Additionally, limitations 1.f and 6.e are not supported by the Provisional. These limitations were added by amendment based on paragraphs 49, 60-61, and 98 of the specification (EX1003, 136), but those paragraphs or similar disclosures are not in the Provisional. Accordingly, it is my opinion that the Provisional does not establish a priority date earlier than October 31, 2006 for the '823 patent.

60. The '823 patent "relates generally to client-server communications and more specifically to causing a browser to store information in a browser storage area of a client device." EX1001, 1:31-34. The information stored in the browser storage area of the client device includes cache cookies, which, unlike a standard cookie, "cannot be blocked or cleared via spyware or a browser setting." EX1001, 5:21-22. The cache cookie is similar to a standard cookie in that it can identify the client device. EX1001, 5:16-18. For example, when "browser 116 establishes a second network session with the server 104 ... the server 104 'reads' the data (the cache cookie 174) ... to identify the client device." EX1001, 5:29-34.

61. The '823 patent discloses two types of cache cookies stored in different areas of the browser storage area. EX1001, Fig. 2. The first type of cache cookie takes the form of URLs stored in the history cache of the browser storage area that the server causes the client device to visit. EX1001, 6:24-27 ("[o]ne way

a server can “write” to the client’s history cache 204 is by redirecting the user to other URLs (within or external to the server’s domain space”)), 6:27-32 (“a server operating the domain www.server.com can redirect a browser to a URL of the form “www.server.com?Z” for any desired value of Z when the browser visits www.server.com, thereby inserting “www.server.com?Z” into the history cache 204 of the client”), 6:39-41 (“the URLs written to the history cache 204 by the server are an embodiment of the cache cookies”).

62. The second type of cache cookie takes the form of temporary internet files (TIF) that are stored in a TIF storage area of the browser. EX1001, 7:1-6 (“[i]n order to place an object X in the TIF area 212, a server can serve content to the browser that causes the browser to download object X” and the “server can verify whether the browser contains object X in its browser storage area 200 by, for example, redirecting the browser to a URL that contains object X”), 7:6-8 (“[i]f TIF X is not present in the browser storage area 200, then the browser requests object X from the server and downloads object X”), 7:8-11 (“[i]f TIF X is present in the TIF area 212 of the browser storage area 200, the browser does not request object X from the server but instead retrieves its local copy”), 7:11-13 (“server can determine whether the browser requests X or retrieves its local copy of X and can use this determination to identify the client device (or user)”).

B. Overview of the '823 Patent's File History

63. During prosecution, the Examiner issued Office Actions rejecting the pending claims under 35 U.S.C. §102 and §103. EX1003, 82-106, 147-171, 200-224. In response, the Applicant amended the claims adding the following limitations:

- “wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the previous network session, wherein the client device initiating the set of network resource request caused the data representative of the set of network resource requests to be stored at the client device.” EX1003, 130-139.
- “perform a first identification of at least one of the client device and a user of the client device using one of the first cookie of the first type and the second cookie of the second type; perform a second identification using a cookie not used in the first identification; and wherein one of the first and second identifications is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device” EX1003, 182-192.
- “wherein the first cookie of the first type is stored in a first client device storage area and the second cookie of the second type is stored in a

second client device storage area different from the first client device storage area ... the cookie not used in the first identification comprising one of the first cookie of the first type and the second cookie of the second type ... perform a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie.” EX1003, 273-282.

64. Following the Notice of Allowance, the Applicant filed an after allowance amendment in which the Applicant amended the claims to include two previous network sessions – “first previous network session” and “second previous network session.” EX1003, 345-350. The Applicant also amended the claims to recite:

- “wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device ... perform a second identification of at least one of the client device and the user of the client device using the second cookie of the second type.” EX1003, 345-350.
- “client device browser storage area.” *Id.*

The Examiner entered the after allowance amendment. EX1003, 356-357.

C. Person of Ordinary Skill in the Art

65. In my opinion, a person of ordinary skill in the art at the time of the alleged invention of the '823 patent (October 31, 2006) (“POSITA”) would have had a Bachelors’ degree in computer science, electrical or computer engineering, or a comparable field of study, plus approximately two to three years of professional experience with computer networks and digital information transmission techniques or other relevant industry experience. Additional graduate education could substitute for professional experience, and significant experience in the field could substitute for formal education.

D. Claim Construction Under 37 C.F.R. §§42.104(b)(3)

66. I understand that in an *inter partes* review, challenged claims are interpreted using the same claim construction standard that is used to construe the claim in a civil action in federal district court. I understand that this means that claim terms of a patent are given the ordinary and customary meaning the terms would have to a POSITA, in view of the description provided in the patent itself and the patent’s file history. I understand that to determine how a POSITA would understand a claim term, one should look to those sources available that show what a POSITA would have understood the disputed claim language to mean. I understand that, in construing a claim term, one looks primarily to the intrinsic patent evidence, including the words of the claims themselves, the remainder of the

patent, and the patent's prosecution history. I understand that extrinsic evidence, which is evidence external to the patent and the prosecution history, may also be useful in interpreting patent claims but that extrinsic evidence cannot be used to vary or contradict the meaning indicated by the intrinsic patent evidence, which I understand is the claims, specification and prosecution history of the '823 patent. I understand that words or terms should be given their ordinary and accepted meaning unless it appears that the inventors were using them to mean something else. In making this determination, the claims, the remainder of the patent, and the prosecution history are of paramount importance. Additionally, the patent and its prosecution history must be consulted to confirm whether the patentee has acted as its own lexicographer (*i.e.*, provided its own special meaning to any disputed terms), or intentionally disclaimed, disavowed, or surrendered any claim scope.

67. I have reviewed Plaintiff's proposed constructions from the prior litigation, and confirmed that Plaintiff has taken the position that no constructions are necessary for any term. EX1011.

68. In my opinion, because the prior art asserted herein discloses the preferred embodiment within the indisputable scope of the claims, there are no terms that need construction. I have been informed that for the purposes of this Declaration, and in accordance with Plaintiff's position from litigation, each of the terms of the Challenged Claims should be accorded their plain and ordinary

meaning as would be understood by a POSITA (person of ordinary skill in the art in question at the time of the invention in April 2006).

E. Patent Owner’s Infringement Contentions

69. I have reviewed what I have been informed is Patent Owner’s current Complaint and Patent Owner’s infringement contention from the prior litigation against Home Depot. EX1007, ¶¶47-74; EX1009.

IV. GROUND 1: HINTON RENDERS OBVIOUS CLAIMS 1-10

A. Overview of Hinton

70. Hinton was filed December 19, 2001, and published as U.S. Publication No. 2003/0115267 on June 19, 2003, and identifies IBM Corporation as its assignee. EX1005, cover page. I have been told to assume that Hinton is prior art under at least 35 U.S.C. §§102(a), 102(b), and 102(e).

71. Hinton generally relates to a system and method for “online user identification, authentication, and authorization” as they relate to “cross-domain log on technologies and technologies which create and manage virtual communities of online users.” EX1005, ¶7. For example, Hinton provides a solution to “cross-domain single-sign-on system and method which allows an Internet user to establish a long-term relationship with participating domains, which gives the user the ability to go directly to participating domains, via bookmarks or direct URLs for example, without having to go through a home

domain first.” EX1005, ¶15. Hinton’s e-community includes a user’s home domain and at least one other domain. EX1005, ¶49, Fig. 1.

72. To implement the cross-domain single-sign-on process, Hinton discloses the use of two cookies: a domain identity cookie (DIDC) and an e-community cookie (eCC). The DIDC is generated when the user enrolls in the e-community and is used to identify the user’s home domain. EX1005, ¶¶70-71. The eCC is a cookie that “acts as an ‘authenticator bookmark’ within a given DNS domain.” EX1005, ¶130. The user has “one e-community cookie set for each domain at which it has a current, authenticated (or vouch-for) session.” EX1005, ¶132. The eCC cookie “indicates the security server or other plug-in location, and a URI at a plug-in location that can provide an authentication ‘vouch for’ token for that user.” EX1005, 133. Accordingly, Hinton discloses the use of two cookies—DIDC and eCC—to facilitate a single-sign-on process across domains.

73. In my opinion, Hinton is from the same field of endeavor as the ’823 patent. They are both patents in the field of computer networks and digital information transmission techniques. In addition, Hinton is reasonably relevant to the problem the patent is concerned with, as shown in my analysis below in section IV.B.

74. Based on my review of the prosecution history (EX1003), Hinton appears to have not been considered during prosecution of the ’823 patent.

B. Motivation to Combine

75. As shown below, no combinations with Hinton are required to arrive at the claimed invention.

C. Claim 1

76. In my opinion, Hinton discloses and renders claim 1 obvious.

1. [1.pre] - A system, comprising:

77. In my opinion, Hinton discloses the preamble.

78. Hinton discloses *a system* (affiliated domain server). EX1005, ¶¶45, 49, Fig. 1, Abstract.

79. Hinton discloses an e-community that includes at least two domains—home domain 103 and an affiliated domain (*e.g.*, other domain 106 and/or another domain 108), each of which is associated with and implemented by one or more servers. EX1005, ¶45 (“an ‘e-community’ has many different ‘participants,’ including e-community members, or domains corresponding to the business units that are participating in the e-community”), ¶49 (“FIG. 1 illustrates a simple e-community architecture, where a user (100) accesses the e-community from their browser” where “[i]n this example, there are three participants in the e-community: the user’s home domain (103), an ‘other’ domain (106) and ‘another’ domain (108)”), (“e-community has, in general, more than two participants”), Abstract (“An Internet user transfers directly to a domain within an e-community without

returning to a home domain or re-authenticating. The user's home domain server prepares and forwards a home domain identity cookie (DIDC) with an enrollment request to a user's browser, with the enrollment request being redirected to an affiliated domain server in the e-community.”).




Figure 1

EX1005, FIG. 1.

80. Thus, in my opinion, Hinton discloses the preamble.

2. [1.a] - one or more processors configured to

81. In my opinion, Hinton discloses and renders obvious this limitation.

82. In my opinion, Hinton discloses and renders obvious *one or more*

processors (processors associated with an affiliated domain server). EX1005,

Abstract, claim 10. Affiliated domain (*e.g.*, other affiliated domain 106 or another affiliated domain 108) is associated with a server. EX1005, Abstract (“user’s home domain server prepares and forwards a home domain identity cookie ... with an enrollment request to a user’s browser, with the enrollment request being redirected to an affiliated domain server in the e-community). Hinton does not explicitly state that the affiliated domain server includes a processor, but a POSITA would have understood and found obvious that a server includes a processor.

3. [1.b.i] – receive a network resource request from a client device wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session

83. In my opinion, Hinton discloses and renders obvious this limitation.

84. In my opinion, Hinton discloses and renders obvious that the affiliated domain’s processor is configured to *receive a network resource request* (access request triggering a vouch-for process) *from a client device* (user’s browser 100) *wherein the network resource request corresponds to a first cookie of a first type* (DIDC). EX1005, ¶¶71, 81, 87-88, 93-97, 137. Hinton’s affiliated domain receives a “vouch-for” request from user 100 to access contents of its domain, which a POSITA would have understood to be a network resource request. EX1005, ¶137 (“vouch-for process occurs when an on-home front-end receives a request from a user that includes a domain identity cookie (DIDC) but not an e-community cookie

(ECC) generated by that front end”). The “vouch-for” request corresponds to a first cookie of the persistent cookie type (*i.e.*, a DIDC cookie) and has a specific format:

```
[0093] DIDC(x)={home domain=103,  
[0094] vouch for URI=www.103.com/101/vouch_  
for.htm,  
[0095] e-community=sample,  
[0096] creation date=Nov 1, 2000,  
[0097] extensible data(x)=data(x)}, hash(info)
```

EX1005, ¶¶93-97. As I explain above in Section I.B., a POSITA understood that a persistent cookie was one known type of cookie. Hinton teaches that the DIDC cookie is persistent. For example, “the user’s browser extracts (62) the DIDC and stores it in the browser’s persistent cookie store, such as string it in a cookie folder on a hard disk drive.” EX1005, ¶88. To the extent that Hinton does not explicitly state that the request is received *from a client device*, in my opinion, a POSITA would have understood that the request would have to come from a client device since the request originates from one, as shown by “user’s browser 100” in Fig. 1. In my opinion, a POSITA would have understood and found it obvious that the affiliated domain’s processor is configured to perform this function.

85. Hinton’s DIDC *was caused to be stored to the client device* (browser’s persistent cookie store) *during a first previous network session* (enrollment session). Hinton’s system includes an enrollment process in which the user’s home domain authenticates the user and generates a DIDC to be stored at the user device

to be used for subsequent network sessions. EX1005, ¶81 (“[f]irst, the user (100) accesses an ‘enroll in e-community’ resource at domain (103), at which time the SSO plug-in (109) at home domain (101) receives (52) this request and checks (53) if the user (100) has authenticated to the home domain (101)”), ¶87 (“plug-in (109) then builds an identity cookie DIDC (103)”), ¶88 (“user’s browser extracts (62) the DIDC and stores it in the browser’s persistent cookie store, such as storing it in a cookie folder on a hard disk drive”). Hinton’s enrollment process is performed before the request triggering the vouch-for process. EX1005, ¶71 (“[t]he purpose of the domain identity cookie is to identify the user’s ‘home’ domain, to identify a URL in the user’s home domain that can ‘vouch for’ the user’s identity, and to identify the e-community in which this user is a participant”). This is further supported by the fact that Hinton’s “vouch-for” process requires the presence of a DIDC. EX1005, ¶137 (“vouch-for process occurs when a non-home front-end receives a request from a user that includes a domain identity cookie (DIDC)”).

4. [1.b.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session

86. In my opinion, Hinton discloses this limitation.

87. Hinton’s *first cookie of the first type (i.e., DIDC) was caused to be stored to the client device at least in part by causing the client device to initiate a*

set of network resource requests (redirections to other domains and back to the home domain) *determined during the first previous network session* (enrollment process). EX1005, ¶¶87-91. As part of the enrollment process, Hinton's home domain builds the DIDC and sends the DIDC to the user device. EX1005, ¶87 ("plug-in (109) ... builds an identity cookie DIDC (103) and an "enrollment token" for the user (100), and creates a response, re-directed to the other community domain (106)'), ¶25 ("[t]he enrollment request is sent via the user's browser using HTTP redirection" and includes "a home domain identity cookie (DIDC) set by the home domain" which the "user's browser extracts and stores the home DIDC"). The home domain's response to the user causes the user to initiate a set of network resource requests by redirecting the response from the home domain to the other domains. EX1005, ¶88 ("user's browser extracts ... DIDC ... and redirects the response to other community domain" and "plug-in (109) at the other domain (106) front-end (104) receives (63) the enrollment request from the home domain front-end (101) which was redirected through the user (100)'), Fig. 6 (element 62), ¶89 ("plug-in (109) at the other domain's front end (104) 'unpacks' the enrollment token, and builds an domain identity cookie for the user for the other domain (106)" which triggers "[a]n 'enrollment successful' message ... sent to the home domain's front end (101) via redirection (63) through the user's browser (100) along with the domain identity cookie for the other domain (106)," which "the

user's browser extracts (64) ... and puts in the browser's persistent cookie store").

The network requests sent from the user to the other domains trigger the other domains to enroll the user at each of the other domains. EX1005, ¶26 (“[a]n enrollment success message is sent by the affiliated domain to the home domain, including the affiliated DIDC and a success indicator. Again, the message is sent via the user's browser using redirection. The user's browser extracts and stores the affiliated DIDC”), ¶90 (“The home domain (103) plug-in (109) at the first front-end (101) receives (65) the redirected ‘enrollment successful at other domain’ message”). As a result of this process, a DIDC, with indicators that the user was enrolled at the other domains, is caused to be stored at the user device. EX1005, ¶90 (“home domain ... modifies (65) the home domain DIDC to include an ‘enrollment success at other domain’ symbol in the extensible attribute data” which “is then returned (65) to the user in the next response from the first front-end (101)”), ¶91 (“[i]n this manner, the home domain DIDC is ‘built up’ or accumulated to include indicators of successful enrollments at affiliated domains within the e-community”). In addition, Hinton teaches that “[i]f the user has not already been authenticated to the home domain, then the SSO plug-in (109) prompts (24) the user (100) for authentication information (e.g. user name and password), and performs (25) authentication verification.” EX1005, ¶108. A POSITA would have understood that this authentication verification requires the

client device to initiate a set of network resource requests determined during the first previous network session to provide such authentication information at the prompting of the SSO plug-in of the home domain.

5. [1.b.iii] - wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device

88. In my opinion, Hinton discloses this limitation. EX1005, ¶¶87-92. During the generation of the DIDC that is stored on the user device based on a client device's network resource requests, the home domain generates an enrollment token for the user which is sent to other domains within the e-community, through a series of network redirects, in order to build-up the home domain DIDC. EX1005, ¶87 (“plug-in (109) ... builds ... an ‘enrollment token’ for the user (100), and creates a response, re-directed to the other community domain (106)”), ¶89 (“plug-in (109) at the other domain’s front end (104) ‘unpacks’ the enrollment token, and builds an domain identity cookie for the user for the other domain (106)” and an “‘enrollment successful’ message is then sent to the home domain’s front-end (101) via redirection (63) through the user’s browser (100) along with the domain identity cookie for the other domain 106” which “the user’s browser extracts (64) ... and puts ... in the browser’s persistent cookie store”), ¶90 (“the home domain (103) plug-in (109) at the first front-end (101) receives (65) the redirected ‘enrollment successful at other domain’ message” and “modifies (65)

the home domain DIDC to include ‘enrollment success at other domain’ symbol in the extensible attribute data” which “is then returned (65) to the user in the next response from the first front-end”), ¶91 (“the home domain DIDC is ‘built up’ or accumulated to include indicators of successful enrollments at affiliated domains within the e-community”), ¶91 (“[t]his process may continue for additional domains in the e-community, using the user’s browser as a re-direction node in the communication path to pass enrollment tokens and success tokens between the home domain and the affiliated domain”). At the end of the network re-directs, on the user’s device “the user’s browser receives a persistent domain identity cookie set by each of the e-community members (103, 106, 108) in which the user has successfully been enrolled.” EX1005, ¶92. This cookie is data representative of the set of network resource requests and it is stored on the client device, and the cookie was received because the client device initiated the set of network resource requests. The home DIDC that is built up to “include the multiple indicators of successful enrollments” constitutes the *data representative of the set of network resource requests that is stored at the client device.*

6. [1.b.iv] - wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session

89. In my opinion, Hinton discloses and renders obvious *a second cookie of a second type (“eCC”) different from the first type was caused to be stored at the*

client device during a second previous network session (limitation 1.b.i). EX1005, ¶35, ¶67, ¶¶93-97, ¶129, ¶130, ¶137, ¶¶147-148, ¶153, ¶¶171-178, ¶249.

90. Hinton also teaches the correspondence between the in-memory e-Community Cookie and a session. “The e-community memory cookies contain security relevant information such that possession of an e-community cookie may provide access to a particular session.” EX1005, ¶232. Similarly, “[a]n eCC is valid for only for the duration of a browser session, and it is expired when a user invokes logout functionality.” EX1005, ¶249.

91. Memory and disk storage as taught by Hinton are two areas in a hierarchy of computer system storage. My own course notes used for teaching Operating Systems in 2004 show storage as organized into a hierarchy with magnetic disk and main store (RAM) as two levels in this storage hierarchy. EX1024 (available at <https://web.cs.wpi.edu/~cs3013/c04/week4-memmgmt.pdf>). Notes from a similar course show a storage/memory hierarchy highlighting speed, cost, size and volatility differences between storage layers. EX1025 (available at <https://www2.cs.uregina.ca/~hamilton/courses/330/notes/memory/MemoryHierarchy.html> with Google search showing a date of Jan 22, 2003).

92. The eCC cookie is a second cookie type (session) different than the DIDC cookie of the persistent cookie type. Hinton discloses an “e-Community Cookie” (eCC) that “is a-memory cookie that is valid within the DNS domain.”

EX1005, ¶130, ¶137, ¶147, ¶153. The DIDC cookie of the persistent cookie type includes the following information:

- [0093] DIDC(x)={home domain=103,
- [0094] vouch for URI=www.103.com/101/vouch_for.htm,
- [0095] e-community=sample,
- [0096] creation date=Nov 1, 2000,
- [0097] extensible data(x)=data(x)}, hash(info)

EX1005, ¶¶93-97 (below row 2, Table 1). In contrast, the eCC cookie is a session cookie type and includes the following information different from the DIDC:

- [0172] eCC(1104)={Auth Server=104,
- [0173] URI at Auth Server=www.106.com/104/vouch_for.htm, e-community=sample,
- [0174] creation date=Nov 1, 2000,
- [0175] extensible attribute=value pairs}, hash-(info)

EX1005, ¶¶172-175 (below row 1, Table 1). As I explain above in Section I.B., a POSITA understood that one type of known cookie was a session cookie type. As shown in these two excerpts, the cookies are not only different types, but also contain different information. In my opinion, a POSITA would have understood and found it obvious that the eCC cookie is a session cookie type that is a different type of cookie from the DIDC cookie of the persistent cookie type.

93. During a session subsequent to the session that created the DIDC, when a user requests access to an affiliated domain, the affiliated domain generates

an eCC for that domain and causes the eCC cookie to be stored at the user's device by sending the eCC cookie to the user. EX1005, ¶137 (“the vouch-for process occurs when a non-home front-end receives a request from a user that includes a domain identity cookie (DIDC) but not an e-community cookie (ECC) generated by that front-end”), ¶177 (“affiliated domain plug-in responds to the user's browser (100) based on the results of the access control decision, including the eCC for the affiliated domain front-end (104)”), ¶178 (“user's browser (100) receives the response and stores the eCC for the affiliated domain's front-end (104) in its cookie store”). In my opinion, Hinton's system discloses saving the eCC cookie in a different session than the DIDC cookie.

7. **[1.b.v] wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area.**

94. In my opinion, Hinton discloses and renders obvious *the first cookie of the first type (DIDC cookie) is stored in a first client device browser storage area and the second cookie of the second type (eCC cookie) is stored in a second client device browser storage area different from the first client device browser storage area*. EX1005, ¶¶88, 233, 245. Hinton discloses that the DIDC cookie is stored in the browser's persistent cookie store while the eCC is stored the browser's cookie memory, which a POSITA would have understood is different

from the browser's persistent cookie store. EX1005, ¶88 (“the user's browser extracts (62) the DIDC and stores it in the browser's persistent cookie store, such as storing it in a cookie folder on a hard disk drive”), ¶233 (“domain identity cookie (“DIDC”) is a persistent cookie that resides in the user's cookie ‘jar’, such as a cookie.txt file”), ¶248 (“the e-community cookie (eCC) ... resides in the user's browser cookie memory”). As Hinton expressly discloses that the storage areas for the two cookies have different names, in my opinion, a POSITA would have understood that the two storage areas are different browser storage areas. In my opinion, a POSITA would have understood and found it obvious that because the browser's cookie store is persistent and the browser cookie memory is temporary, they are different storage areas for the browser.

8. **[1.c] - based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determine information that was encoded and stored in the client device;**

95. In my opinion, Hinton discloses and renders obvious that the affiliated domain's processor is configured to, *based at least in part of the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session* (limitation 1.b.i), *determine information* (user's home domain information) *that*

was encoded and stored in the client device. EX1005, ¶¶137, 143, 155, 236, 239, 245.

96. Hinton’s affiliated domain *determines information* from the DIDC cookie (first cookie stored in previous network session, limitation 1.b.i) in the request. Specifically, when the affiliated domain receives the request from the user to access a resource protected by a plug-in at the affiliated domain, the plug-in “‘parses’ the DIDC(106) [sic] cookie to determine the user’s home domain, a URI in their home domain that can vouch for the user’s identity, the e-community in which they are enrolled, and a creation/update timestamp.” EX1005, ¶155, ¶236 (“vouch for resource will be a URL that contains some form of active content that can authenticate the user ... and build a ‘vouch for’ token”). When the affiliated domain receives a network resource request from a client device that includes a DIDC, the DIDC corresponds to a DIDC created on the client device in a prior network session. In my opinion, a POSITA would have understood and found it obvious that Hinton discloses the affiliated domain’s processor is configured to perform this function.

97. The information in the DIDC (stored in the client device) is hashed, which is a type of encoding of stored data. EX1005, ¶239 (“[t]he information is hashed for integrity protection”), ¶245 (“the domain identity cookie may or may not be protected by keyed hash” which “will at most provide integrity protection

on the data in the cookie”). In my opinion, a POSITA would have understood and found obvious that hashing the information in the DIDC, such as through a keyed hash, constituted an encoding of the information.

9. [1.d.i] - perform a first identification of at least one of the client device and a user of the client device using the first cookie of the first type,

98. In my opinion, Hinton discloses and renders obvious that the affiliated domain’s processor is configured to *perform a first identification of at least one of the client device and a user of the client device using the first cookie of the first type*. EX1005, ¶¶92, 137-139, 141-144, 161-162.

99. In my opinion, a POSITA would have understood that the affiliated domain’s processor identifies the client device using the DIDC, which is stored in a user’s browser’s persistent cookie store on a user’s device. EX1005, ¶92 (“the user’s browser receives a persistent domain identity cookie set by each of the e-community members ... in which the user has successfully been enrolled”), ¶27 (“updated home DIDC is then transmitted to the user’s browser, where it is stored in the persistent cookie store”), ¶¶88-89. The affiliated domain uses the DIDC cookie in subsequent sessions to identify that the requesting device has been previously authorized. EX1005, ¶143 (“[a] prerequisite for the transfer of authentication information across domains is that the user has already enrolled in the e-community” and “[i]f there is no DIDC cookie, the front-end will treat the

user as a ‘normal’ internal user (as opposed to a participant in the e-community) and will attempt to authenticate the user”). In this manner, Hinton’s DIDC is used to identify the client device. As a DIDC cookie is stored in a client device’s browser, in my opinion, a POSITA would have understood and found it obvious that the use of the DIDC cookie is the performance of an identification of the client device. EX1005, ¶27. In my opinion, a POSITA would have understood and found it obvious that Hinton discloses the affiliated domain’s processor is configured to perform this function.

100. The affiliated domain’s processor also identifies the user using the DIDC. EX1005, ¶71 (“[t]he purpose of the domain identity cookie is to identify the user’s ‘home’ domain, to identify a URL in the user’s home domain that can ‘vouch for’ the user’s identity, and to identify the e-community in which this user is a participant”), ¶137 (“vouch-for process occurs when a non-home front-end receives a request from a user that includes a domain identity cookie (DIDC) but not an e-community cookie (ECC) generated by that front end”). Using the DIDC, the affiliated domain is able to identify the user by requesting a vouch-for token from the user’s home domain. EX1005, ¶139 (the vouch for process includes “[r]equesting ... the user’s home domain to ‘vouch for’ the user”), ¶141 (the vouch for process includes “[g]eneration of a ‘vouch for token’ (VT) to transfer back to the requesting domain a redirected response”), ¶144 (“VT is used to vouch for the

authenticity of the user's identity to the other e-community domains"). The vouch for token includes an indication of the identity of the user in the form of a userID. EX1005, ¶¶161-162 ("first front-end (101) builds a 'Vouch-For Token' (VT) to provide the vouch-for information to the affiliated domain (106) such as: VT=E{Tag=VT, userid=jsmith, homedomain=103"). Thus, because the DIDC is used to generate a vouch-for token, and because the vouch-for token includes an identification of the user, the DIDC is used to perform a first identification of the user.

101. In my opinion, a POSITA would have understood and found it obvious that Hinton discloses the affiliated domain's processor is configured to perform this function.

10. [1.d.ii] wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device

102. In my opinion, Hinton discloses and renders obvious this limitation.

103. Hinton discloses that *the first identification is performed using the first cookie of the first type*. §IV(C)(9) (limitation 1.d.i). The information in the DIDC is encoded and stored in the client device. *See* §IV(C)(3)-(4), (B)(8) (limitations 1.b.i-1.b.ii, 1.c).

104. Hinton further discloses that the first identification is performed *in part by using the determined information that was encoded and stored in the client*

device (vouch for URI). EX1005, ¶¶138-139; §IV(C)(8) (limitation 1.c) (showing vouch for URI is determined information from the DIDC cookie). The DIDC (*i.e.*, first cookie) includes information, such as the vouch for URI, which is used by the affiliated domain for the vouch-for process. *See* §IV(C)(3) (limitation 1.b.i) (DIDC includes “vouch for URI=www.103.com/101/vouch_for.htm”); EX1005, ¶¶138-139 (the vouch for process includes “(1) Identification that user is in the e-community but has a different home domain; (2) Requesting (via re-direction) the user’s home domain to ‘vouch for’ the user”). Hinton discloses the DIDC’s data is hashed (EX1005, ¶245), a POSITA would have understood the hashing to include the identification information. Thus, the first identification is performed at least in part using the determined information that was encoded and stored in the client device. In my opinion, a POSITA would have understood and found it obvious that Hinton discloses the affiliated domain’s processor is configured to perform this function.

11. [1.e] - perform a second identification of at least one of the client device and the user of the client device using the second cookie of the second type

105. In my opinion, Hinton discloses and renders obvious that the affiliated domain’s processor *performs a second identification of at least one of the client device and the user of the client device using the second cookie of the second type* (eCC cookie). EX1005, ¶¶129, 132, 133, 154, 245.

106. When a user requests to access resources in an affiliated domain, the affiliated domain creates an eCC for the affiliated domain front-end once the user is authenticated and vouched for. EX1005, ¶130 (“[a]s a result of authentication, the SSO plug-in generates an ‘e-Community Cookie’ (an eCC or e-community cookie) that acts as an ‘authenticator bookmark’”), ¶132-33 (“a user has one e-community cookie set for each domain at which it has a current, authenticated (or vouched-for) session”). Once the eCC is created for the affiliated domain, the user can use the eCC to access any server in the domain. EX1005, ¶133 (“[o]nly the one instance with a DNS domain that authenticates the user or first receives an authentication ‘vouch-for’ message sets an e-community cookie at the user’s browser” and the eCC “is a domain cookie and can therefore be sent to any server in the domain that created it” which “allows for simplified single-sign-on capabilities within a domain that is partitioned by multiple security server domains”). When the user requests access to one of the servers in the domain that created the eCC, that server checks whether the user has an eCC to access the DNS domain it is a part of. EX1005, ¶154 (“plug-in at the associated domain front-end (104) looks for an eCC cookie set by a different front-end within the associated domain (106)” and “[i]f present, this would indicate that the user has a session with a different front-end within the associated domain (106)”). In my opinion, a

POSITA would have understood and found it obvious that Hinton's affiliated domain processor is configured to perform this function.

107. The eCC is used to identify the user of the client device. EX1005, ¶245 (“[t]he eCC is used to identify the Web server cluster, within a given domain, that can vouch for a user’s identity”), ¶131 (“eCC identifies the server that authenticated the user and a URI pointing to an authentication script that can vouch for the user within a given domain”). In my opinion, a POSITA would have understood and found it obvious that Hinton's affiliated domain processor is configured to perform this function.

108. In my opinion, a POSITA would have understood that the affiliated domain's processor identifies the client device using the eCC, which is stored in the browser's cookie memory on the user's device. *See* §IV(C)(7) (limitation 1.b.v). Hinton discloses the eCC cookie is stored in a client device's browser's cookie memory. EX1005, ¶167, ¶248. In my opinion, a POSITA would have understood and found it obvious that the use of the eCC cookie is the performance of an identification of the client device. EX1005, ¶167, 248.

109. In my opinion, a POSITA would have understood and found it obvious that Hinton discloses the affiliated domain's processor is configured to perform this function.

- 12. [1.f] - perform a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie**

110. In my opinion, Hinton discloses and renders obvious that the affiliated domain's processor is configured to *perform a determination based at least in part of (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie.* EX1005, ¶137. Hinton discloses that the DIDC is generated and stored to the client device based on a network resource request. *See* §IV(C)(4) (limitation 1.b.ii). Hinton discloses that an eCC is generated and stored to the client device based on a network resource request. *See* §IV(C)(6) (limitation 1.b.iv). Hinton further discloses that the vouch-for process is performed based on the presence of the DIDC and the absence of the eCC. EX1005, ¶137 (“vouch-for process occurs when a non-home front-end receives a request from a user that includes a domain identity cookie (DIDC) but not an e-community cookie (ECC) generated by that front-end”). Thus, Hinton's vouch-for process makes a determination based on the presence of a network request associated with the first cookie (*i.e.*, the network request that triggered the vouch-for process, which included the DIDC) and associated with an absence of a second cookie (*i.e.*, the user does not have an eCC associated with the affiliated

domain). In my opinion, a POSITA would have understood and found it obvious that the affiliated domain's processor is configured to perform this function.

13. [1.g] - a memory coupled to the one or more processors and configured to provide the one or more processors with instructions.

111. Hinton discloses *a memory (e.g., RAM) coupled to the one or more processors and configured to provide the one or more processors with instructions.* EX1005, Abstract, claim 10; §IV(C)(2) (limitation 1.a). The affiliated domain (*e.g.*, other affiliated domain 106, another affiliated domain 108) includes a server. EX1005, Abstract (“user’s home domain server prepares and forwards a home domain identity cookie ... with an enrollment request to a user’s browser, with the enrollment request being redirected to an affiliated domain server in the e-community). In my opinion, a POSITA would have understood and found it obvious for the affiliated domain’s server to include a memory coupled to its processor and configured to provide the processor with instructions. *See e.g.*, EX1005, Claim 10 (“A computer readable medium encoded with software for allowing an Internet or intranet browser user to transfer directly to a domain that is participating in an e-community without repetitious and redundant authentication actions, said e-community comprising a plurality of affiliated domain servers, said user being properly registered and authenticated to a home domain server within said e-community, said software causing a processor to perform the steps ...”).

D. Claim 2 - The system of claim 1 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.

112. In my opinion, Hinton discloses and renders this claim obvious.

113. As I explain above in Section IV(C), Hinton discloses and renders obvious *the system of claim 1*. See §IV(C) (claim 1).

114. Hinton also discloses that “a user will access resources in different (‘participating’) domains on behalf of their home domain.” EX1005, ¶11. For example, “the home domain itself may have ‘long term’ relationships with other domains.” EX1005, ¶10. Hinton’s invention addresses the “need in the art for a cross-domain single-sign-on system and method which allows an Internet user to establish a long-term relationship with participating domains, and which gives the user the ability to go directly to participating domains, via bookmarks or direct URL’s.” EX1005, ¶15. In the case where a user requests a network resource from an affiliated domain (*i.e.*, not the home domain that generates the DIDC and the eCC), Hinton discloses that the affiliated domain identifies the user. EX1005, ¶147 (“[v]ouch for information is transferred across domains when a user requests a resource in a domain other than their home domain, where the request requires an authenticated identity [of the user]”). EX1005, ¶¶84-85, ¶143, ¶45.

E. Claim 3 - The system of claim 1 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.

115. In my opinion, Hinton discloses and renders this claim obvious.

116. As I explain above in Section IV(C), Hinton discloses and renders obvious *the system of claim 1*. See §IV(C) (claim 1).

117. The client device is identifiable by a first cookie in a first client device browser storage area. §IV(C)(7), (9) (limitations 1.b.v, 1.d.i), The client device is also identifiable by a second cookie in a second client device browser storage area. §IV(C)(7), (11) (limitations 1.b.v, 1.e).

F. Claim 4 - The system of claim 1 wherein the determination comprises detection of pharming.

118. In my opinion, Hinton discloses and renders this claim obvious.

119. As I explain above in Section IV(C), Hinton discloses and renders obvious *the system of claim 1*. See §IV(C) (claim 1).

120. Hinton's DIDC includes various information about the servers in the e-community, such as the URI corresponding to the home domain for purposes of initiating the vouch-for process. EX1005, ¶¶93-97. Hinton's eCC also includes the URI corresponding to the affiliated domain that issued the eCC. A POSITA would have been familiar with the problem of domain spoofing, or pharming ((EX1015) "The Pharming Guide (part 2)" Dec 14, 2004 available at:

<http://www.technicalinfo.net/papers/Pharming2.html>). And, in my opinion, a POSITA would have been motivated to add detection of pharming to Hinton. In order to prevent possible domain spoofing, a POSITA would have been motivated by server load balancing to encode a server identifier, such as an Internet Protocol (IP) address, for the home domain as an additional field in the DIDC and/or a server identifier, such as an IP address, of the affiliated domain as an additional field in the eCC. *See e.g.*, EX1013. In my opinion, a POSITA would have understood and found obvious that the affiliated domain's processor would use the home domain server's IP address and/or the affiliated domain server's IP address as an additional check to determine whether a malicious actor spoofed the home domain server's URI or the affiliated domain server's URI, respectively, *i.e.*, perform a detection of pharming, based on the received DIDC or eCC that includes the URI and the IP address of each respective server.

G. Claim 5 - The system of claim 1 wherein in response to the performed determination, the one or more processors are configured to cause a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.

121. In my opinion, Hinton discloses and renders this claim obvious.

122. As I explain above in Section IV(C), Hinton discloses and renders obvious *the system of claim 1*. *See* §IV(C) (claim 1).

123. In the case where a user requests a network resource from an affiliated domain for which an eCC was generated and then deleted (*i.e.*, the DIDC is present, but not the eCC), the affiliated domain’s processor generates and sends to the client device a new eCC corresponding to its domain, which is a third cookie of the second type. EX1005, ¶45 (“an ‘e-community’ has many different ‘participants’, including e-community members, or domains corresponding to the business units that are participating in the e-community”), ¶143 (“[a] prerequisite for the transfer of authentication information across domains is that the user has already enrolled in the e-community”), ¶147 (“[v]ouch for information is transferred across domains when a user requests a resource in a domain other than their home domain, where the request requires an authenticated identity”), ¶249 (“[a]n eCC is valid for only for the duration of a browser session, and it is expired when a user invokes logout functionality”). Thus, the replacement token is a third cookie of one of the second type.

H. Claim 6

124. In my opinion, Hinton discloses and renders claim 6 obvious. *See* §IV(C) (claim 1).

1. [6.pre] A method

125. In my opinion, Hinton discloses and renders obvious the preamble.

Above in Section §IV(C), I show that Hinton discloses a system that performs a method. *See* §IV(C) (claim1).

2. [6.a.i] receiving a network resource request from a client device, wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session,

126. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(3) (limitation 1.b.i).

3. [6.a.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session,

127. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(4) (limitation 1.b.ii).

4. [6.a.iii] wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device,

128. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(5) (limitation 1.b.iii).

5. **[6.a.iv] wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session,**

129. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(6) (limitation 1.b.iv).

6. **[6.a.v] and wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area**

130. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(7) (limitation 1.b.v).

7. **[6.b] based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determining information that was encoded and stored in the client device; and**

131. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(8) (limitation 1.c).

8. **[6.c.i] performing a first identification of at least one of the client device and a user of the client device using the first cookie of the first type**

132. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(9) (limitation 1.d.i).

- 9. [6.c.ii] wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device; and**

133. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(10) (limitation 1.d.ii).

- 10. [6.d] performing a second identification of at least one of the client device and the user of the client device using the second cookie of the second type; and**

134. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(11) (limitation 1.e).

- 11. [6.e] performing a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request**

135. In my opinion, Hinton discloses and renders obvious this limitation.

See §IV(C)(12) (limitation 1.f).

- I. Claim 7 - The method of claim 6 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.**

136. In my opinion, Hinton discloses and renders obvious this claim. *See*

§§IV(D), IV(H) (claims 2, 6).

J. Claim 8 - The method of claim 6 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.

137. In my opinion, Hinton discloses and renders obvious this claim. *See* §§IV(E), IV(H) (claims 3, 6).

K. Claim 9 - The method of claim 6 wherein the determination comprises detection of pharming.

138. In my opinion, Hinton discloses and renders obvious this claim. *See* §§IV(F), IV(H) (claims 4, 6).

L. Claim 10 - The method of claim 6 wherein in response to the performed determination, further comprising causing a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.

139. In my opinion, Hinton discloses and renders obvious this claim. *See* §§IV(G)-(H) (claims 5-6).

140. To the extent it is asserted that the portions of Hinton that I have cited above relate to different, incompatible embodiments (which they do not), in my opinion, it would have been obvious to a POSITA to combine such embodiments into a single system for client-server communications at least because such embodiments are described in the same prior art reference, are fully compatible with each other, and could be combined with minimal effort to achieve predictable results.

141. To the extent it is asserted that Hinton does not disclose any limitation of the challenged claims, in my opinion, such limitation would have nonetheless been obvious to a POSITA in light of Hinton and a POSITA's knowledge. Practicing any limitation of the challenged claims in light of Hinton would have been within the knowledge and skill of a POSITA, would have required minimal effort, would have yielded predictable results, would have been fully compatible with the Hinton system, and would have been a mere design choice. Motivation to do so arises from at least common sense and the disclosures of Hinton that I have set forth above.

V. GROUND 2: VARGHESE RENDERS OBVIOUS CLAIMS 1, 3-6, AND 8-10

A. Overview of Varghese

142. Varghese was filed April 28, 2006, and issued as U.S. Patent No. 7,908,645 on March 15, 2011. EX1004, cover page. I have been told to assume that Varghese is prior art under 35 U.S.C. §102(e).

143. Varghese is generally directed to “systems and methods for providing protection against identify theft of a computer network.” EX1004, 1:16-18. To protect computer networks from theft, Varghese discloses the use of two cookies—secure cookies and flash cookies. EX1004, 5:64-66 (“[t]he present invention includes secure cookies, flash objects and other technologies to recognize and to fingerprint [] from which device a user access[sic] an application”), 6:2-7

("[i]nformation concerning these user devices is fingerprinted and stored into a device token or device id for one-time use," which "is stored on the user device and saved in a database for later comparison with tokens retrieved from subsequent user device accesses"), 6:8-14 ("[t]he present invention also includes user device tokens or device ids that have a unique number which is randomly generated by methods of this invention," which "are ... assigned to the particular user device, stored on the particular user device as persistent data (e.g., a cookie), and also stored so as to be accessible to the authentication services of this invention").

144. The flash cookie and/or the secure cookie are used to identify the user device on subsequent login or authentication attempts. EX1004, 24:23-27 ("[u]ser devices are preferably identified using secure cookies, Flash cookies, and similar data tokens combined with other data items such as browser characteristics, device hardware configuration (e.g., as acquired using Flash calls), network characteristics, and the like"). Secure cookies are standard cookies secured against modification or tampering. EX1004, 25:25-32 ("[a] standard cookie is a data packet sent by a web server to a web browser for saving to a file on the host machine" and "[a] secure cookie refers to a standard cookie that has been secured against modification or tampering"). Flash cookies differ from secure cookies in that they are not as easily removed from the user's device because they are generated and stored by the Flash software on the user device. EX1004, 25:37-43

("[Flash] software can create local shared objects, known as 'flash cookies', for maintaining locally persistent data on a user's device akin to the standard 'cookies' stored by web browser" and "have the advantage not being as easily removed from the user's device as are standard cookies"). Varghese's server uses both the flash cookie and the secure cookie to verify the user. For example, Varghese discloses in Table 8 the various scenarios based on the presence or absence of either of the secure cookie or the flash cookie. EX1004, 19:54- 65 ("This table returns a score of '0' (a score indicated a low likelihood of fraud) in case all evaluated data items are present and match in connection with a current user request. If no data item is present or if all data items do not match, a score of '10' (a score indicated a high likelihood of fraud) is returned. In case where some data items are present and match while other data items are absent or do not match, this table invokes further checks. If the retrieved data tokens that were previously stored on a device by this invention, e.g., a secure cookie, a Flash cookie, or Flash data, are not present, a further pattern check is performed.").

TABLE 8

Primary device decision table					
Data item					
Secure cookie	Flash cookie	Flash data	Browser characteristics	Operating system characteristics	Score
*	*	*	*	*	0
X	*	*	*	*	PATTERN CHECK
M	*	*	*	*	SECONDARY CHECK
X/M	X	*	*	*	PATTERN CHECK
X/M	M	*	*	*	SECONDARY CHECK
X/M	X/M	X	*	*	PATTERN CHECK
X/M	X/M	M	*	*	SECONDARY CHECK
X/M	X/M	X/M	M	*	SECONDARY CHECK
X/M	X/M	X/M	X/M	M	10

Key:
 X = missing;
 M = present and mismatched;
 * = present and matched

EX1004, Table 8 [annotation added].

145. The '823 specification recognizes that a browser storage area for cookie types on a client does not need to be as a result of “an explicit browser feature, but rather a form of persistent state in the browser 116 that the server 104 can access.” EX1001, 5:14-16. Flash, which caches its Local Shared Objects (Flash cookies) in storage accessible to the browser, was known to be an available browser plugin such as described in EX1027 (“HOWTO: Installing Flash Plugin in Firefox Way” by Thomas Chung on Sep 16, 2004

(<https://fedoranews.org/tchung/firefox-flash/>) and shown as part of the browser in EX1028 (a White Paper “Macromedia® Flash® Platform Security and Macromedia Enterprise Solutions” by Adrian Ludwig in September 2005 (https://www.adobe.com/platform/whitepapers/flashplatform_security_enterprise.pdf)).

146. In my opinion, Varghese is from the same field of endeavor as the ’823 patent. They are both patents in the field of computer networks and digital information transmission techniques. In addition, Varghese is reasonably relevant to the problem the patent is concerned with, as shown in my analysis below in Section V.B.

147. Based on my review of the prosecution history (EX1003), Varghese appears to have not been considered during prosecution of the ’823 patent.

B. Motivation to Combine

148. As shown below, no combinations with Varghese are required to arrive at the claimed invention.

C. Claim 1

149. As shown below, Varghese renders claim 1 obvious.

1. [1.pre] - A system, comprising:

150. In my opinion, Varghese discloses the preamble.

151. Varghese discloses *a system* (collectively, system 1302, system 1304, and authentication server 1306). EX1004, 8:40-44 (“FIG. 13A illustrates an exemplary embodiment of the present invention directed to providing authentication services to online service providers who make available to individual users online server applications” that “execute on service-provider machines”), 8:46-48 (“Fig. 13A illustrates ... one or more service-provider computer systems, e.g., systems 1302 and 1304 and an authentication system server system 1306, interconnected through network 710 to one or more user devices 720”), 8:60-9:3 (“In many preferred embodiments ... authentication processes of the invention are implemented with a client-server-type architecture (or, more generally, a distributed-systems-type architecture.”), Fig. 13A.

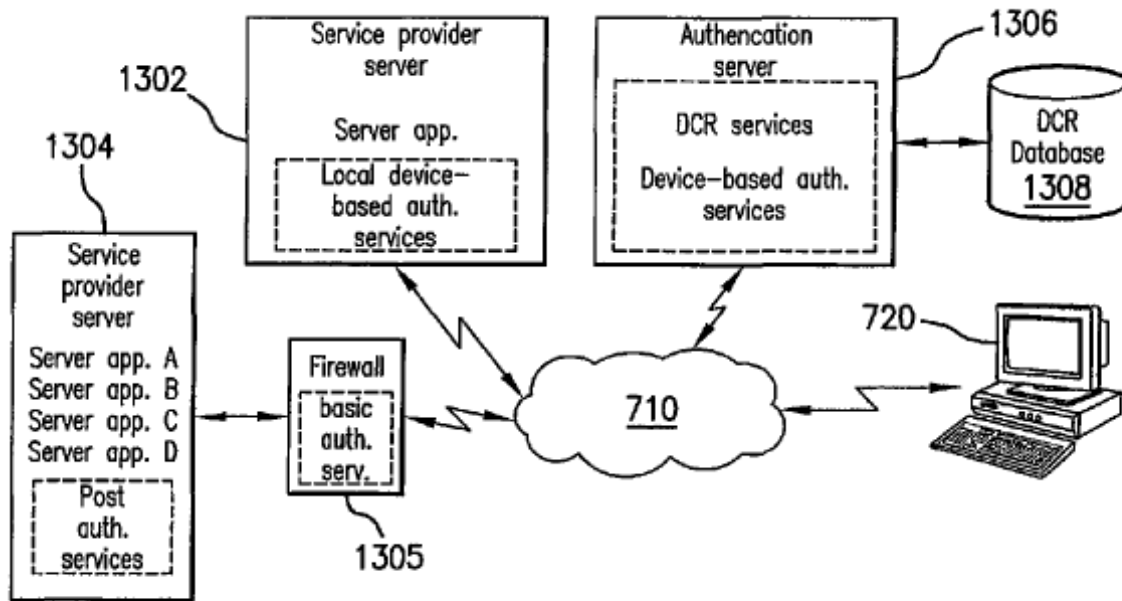


FIG. 13A

EX1004, Fig. 13A.

2. [1.a] - one or more processors configured to

152. In my opinion, Varghese discloses and renders obvious this limitation.

153. Varghese discloses and renders obvious *one or more processors* (processor associated with authentication server 1306). Varghese discloses that authentication server 1306 “is generally structured as known in the art, and includes a CPU.” EX1004, 8:50-51.

3. [1.b.i] – receive a network resource request from a client device

154. In my opinion, Varghese discloses and renders obvious this limitation.

155. Varghese discloses and renders obvious the authentication server processor configured to *receive a network resource request from a client device*. Varghese discloses that server 1306 receives network resource requests from a client device. EX1004, Fig. 13A, 6:21-23 (“[t]he present invention enables application service providers score risk for each online login and transaction and to increase authentication security in real time”), 10:17-19 (“[a]uthentication services are invoked when a server application or services provider computer system receives user request 1320 that needs authentication”), 10:20-21 (“the most common user request is a login request to access an application or system”), 24:40-41 (“request is received at a service provider server from a user device 720.”), 10:20-23, (“user request is a login request to access an application or

system ... transaction requests.”). In my opinion, a POSITA would have understood and found it obvious that authentication server’s processor is configured to receive those requests.

4. **[1.b.ii] - wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session, wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session**

156. In my opinion, Varghese discloses and renders obvious this limitation.

157. Varghese discloses the *network resource request* (authentication request) *corresponds to a first cookie of a first type* (secure cookie). For example, when a user in Varghese attempts to authenticate themselves, their browser provides the authentication server with the secure cookie that was generated during a previous session with the server. EX1004, 24:40-41 (“[i]n Step 402, a request is received at a service provider server from a user device 720”), 15:31-34 (“[a] further important component of device information when available is a secure token, e.g., a secure cookie, available from a device which has been previously used as a user device”), 15:34-40 (“[w]hen a request is received from and device, at least the available location and device information can be summarized, condensed, or fingerprinted and stored back on the device as a secure token” and “[i]f another request then originates from this device, the secure token can be

retrieved and its contents compared against the currently-collected location and device information.”), 2:4-8, (“[a] cookie generally refers to a packet of information, often sensitive information, sent by a web server to a browser resident on the user’s computer system for saving to a file and for transmitting back to the server whenever the user’s browser makes additional requests from the server”).

158. Varghese discloses that the first cookie of the first type *was caused to be stored to the client device*. Varghese discloses that the client device receives and stores the secure cookie that was generated by the server. EX1004, 25:25-29 (“[a] standard cookie is a data packet sent by a web server to a web browser for saving to a file on the host machine,” which “can be retrieved and submitted back to the server when requested”), 25:25-32, 26:4-6 (“a new Device ID token is created for the device ... and ... sent to the user device and stored thereon, e.g., as a standard cookie or as a flash cookie”), 26:6-11 (“[i]f no Device ID was found on the user device, a new Device ID token is created from the gathered identifying information” and “[i]f a Device ID was found, it can be updated, e.g., with a new unique bit string, new timestamp, and so forth”), 24:67-25:3 (“[s]ome or all of the device identity (along with identifying information generated during the fingerprinting process) information is stored in a data token referred to as a ‘Device ID’”).

159. Varghese further discloses that the first cookie was caused to be stored to the client device *at least in part by causing the client device to initiate a set of network resource requests*. Varghese discloses that the secure token was generated during a login in attempt by the user. EX1004, 24:40-42 (“a request is received at a service provider server from a user device 720 ... for data resident thereon”), 26:13-14 (“[a] feature of the invention relates to the replacement of the cookie on the user’s machine upon each login”), 10:17-19 (“[a]uthentication services are invoked when a server application or services provider computer system receives user request 1320 that needs authentication”). In my opinion, a POSITA would have understood that a login prompt is a network request that the server causes the user to initiate a set of network resource requests.

160. Varghese teaches details of authentication along with illustrating user interfaces for doing so with a preferred embodiment making use of Flash software being sent to the user device, which a POSITA would recognize as requiring a subsequent request for a Flash object (i.e., a network resource request):

“Database 704 can include user interfaces such as the interface 18 shown in FIG. 1 and the interfaces shown in FIGS. 2, 3, 8 and 9 and variations thereof, e.g., as shown in FIGS. 10A-E, that latter all illustrate a plurality of higher security interfaces based on the keypad user interface in FIG. 2. The user interfaces are shown with two authentication factors, userid or username and password, it should be appreciated that the present invention is not limited to two factors; additional factors may be included within the scope of the present invention. Each of the GUI’s shown are sent to the user device using suitable software, *e.g.*, MACROMEDIA Flash, Java, etc. In preferred embodiments, Flash software is used for the GUIs. The GUIs in

FIGS. 8 and 9 provide heightened security by enabling users to enter and Submit passwords or other sensitive information using an online image that is unrecognizable in actual image form as well as in data output form.”

EX1004, 29:42-58.

161. That set of network requests were *determined during the first previous network session*. As provided above, the secure token is generated upon an authentication request and then used by the user in subsequent network sessions. *See e.g., EX1004, 15:31-34* (“[a] further important component of device information when available is a secure token, e.g., a secure cookie, available from a device which has been previously used as a user device”). Thus, Varghese discloses that the secured cookie was generated and stored during a network session that preceded the network resource request.

5. [1.b.iii] - wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device

162. In my opinion, Varghese discloses this limitation.

163. Varghese further discloses *the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device*. Varghese discloses that the secure cookie includes information representative of the login request. EX1004, 10:53-60 (“fingerprinting process then gathers identifying information describing the device from which the user request originated and creates a device identifier ‘Device ID’”

that “is stored on the user device from which it can be retrieved and form part of the device identifying information to be used during a subsequent fingerprinting”), 24:50-51 (during the fingerprinting process “device identity information for the user device is captured”), 24:67-25:3 (“[s]ome or all of the device identity (along with identifying information generated by the fingerprinting process) information is stored in a data token referred to as a (‘Device ID’)”). The client device caused this information to be stored at the client device when the client device received the secured cookie for storage. EX1004, 6:2-4 (“[i]nformation concerning these user devices is fingerprinted and stored into a device token or device id for one-time use”), 6:4-7 (“[t]he id or token is stored on the user device and saved in a database for later comparison with tokens retrieved from subsequent user device accesses”).

6. [1.b.iv] - wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session

164. In my opinion, Varghese discloses and renders obvious this limitation.

165. Varghese discloses *a second cookie of a second type different from the first type*. Varghese discloses use of a flash token. EX1004, 25:33-40. Flash cookies are locally shared objects created by flash software and are thus of a second type different from the first type “secure” cookies. EX1004, 25:40-43 (“[f]lash cookies can be stored locally on a flash plug-in user’s device, are

updatable, and have the advantage not being as easily removed from the user's device as are standard cookies"), EX1004, 25:25-32.

166. Varghese discloses that the flash cookie *was caused to be stored at the client device during a second previous network session*. A POSITA would recognize that Flash software sent to the user device for the authentication interface would allow the software to store Flash cookies to be stored at the client device during a second previous network session "Each of the GUI's shown are sent to the user device using suitable software, e.g., MACROMEDIA Flash, Java, etc. In preferred embodiments, Flash software is used for the GUIs." EX1004, 29:51-54. The Flash software allowing the software to store the cookie was one of the events that *caused* a cookie to be stored. Varghese further discloses that secure cookies can be removed from the browser's memory when the user clears the browser's cookies. EX1004, 25:40-43 ("[f]lash cookies ... have the advantage not being as easily removed from the user's device as are standard cookies"). Varghese further discloses that the cookies are routinely replaced with each login (*i.e.*, new network session). EX1004, 26:13-14 ("[a] feature of the invention relates to the replacement of the cookie on the user's machine upon each login"). Thus, in my opinion, a POSITA would have understood and found obvious that Varghese discloses a scenario in which the flash cookie was created and stored in a different (*i.e.*, second) previous network session than the secure cookie (*e.g.*, the user

cleared the browser's cookies). Further Varghese contemplates the cookies being created in different network sessions when only one of the secured cookie or the flash cookie is present. *See, e.g.,* EX1004, Table 8 (showing various scenarios in which one of secure cookie present or flash cookie is present).

TABLE 8

Primary device decision table					
Data item					
Secure cookie	Flash cookie	Flash data	Browser characteristics	Operating system characteristics	Score
*	*	*	*	*	0
X	*	*	*	*	PATTERN CHECK
M	*	*	*	*	SECONDARY CHECK
X/M	X	*	*	*	PATTERN CHECK
X/M	M	*	*	*	SECONDARY CHECK
X/M	X/M	X	*	*	PATTERN CHECK
X/M	X/M	M	*	*	SECONDARY CHECK
X/M	X/M	X/M	M	*	SECONDARY CHECK
X/M	X/M	X/M	X/M	M	10

Key:
 X = missing;
 M = present and mismatched;
 * = present and matched

7. **[1.b.v] - wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area.**

167. In my opinion, Varghese discloses and renders obvious this limitation.

168. Varghese discloses and renders obvious *the first cookie of the first type (secure cookie) is stored in a first client device browser storage area and the second cookie of the second type (flash cookie) is stored in a second client device browser storage area different from the first client device browser storage area.* In my opinion, a POSITA would have known that the secure cookie is stored in a first client device browser storage area different from the first client device browser storage area where the flash cookie is stored. Secure cookies are standard cookies known to be saved by the Web browser in browser-maintained files. EX1004, 25:25-27 (“A standard cookie is a data packet sent by a web server to a web browser for saving to a file on the host machine.”), EX1004, 25:25-32. In contrast, Flash cookies are stored in a separate storage area. EX1004, 25:40-43 (“Flash cookies can be stored locally on a flash plug-in user’s device, are updatable, and have the advantage not being as easily removed from the user’s device as are standard cookies.”). Similarly, the previously cited EPIC article (EX1014), “Local Shared Objects—‘Flash Cookies,’” provides explicit information for where browser-accessible Flash cookies are stored on the client machine as Local Shared Objects: “Windows C:\Documents and Settings\[username]\Application Data\Macromedia\Flesh Player; Macintosh OSX /Users/[username]/Library/Preferences/Macromedia/Flesh Player; GNU-Linux ~/.macromedia.” EX1014; *see also* §IV(C)(7) (Hinton limitation 1.b.v).

8. **[1.c] - based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determine information that was encoded and stored in the client device;**

169. In my opinion, Varghese discloses and renders obvious this limitation.

170. Varghese discloses and renders obvious the authentication server's processor configured to, *based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session (see §V(C)(3)-(4) (limitations 1.b.i-1.b.ii), determine information that was encoded and stored in the client device.* Varghese discloses that the "captured device information," such as the Device ID that is stored as a secure cookie includes persistent data. EX1004, 6:8-14 ("[t]he present invention also includes user device tokens or device ids that have a unique number which is randomly generated by the methods of this invention" and "are then assigned to the particular user device, stored on the particular user device as persistent data (e.g., a cookie), and also stored so as to be accessible to the authentication services of this invention"). Varghese further discloses that the persistent data includes information that are encrypted to secure against modification. EX1004, 25:10-15 ("[s]ecure persistent data includes generally data elements that are encrypted, signed, or otherwise secured against modification, and remain resident on the user device even when it is not accessing

a service provider application”). This persistent data is used by the server to identify the user on subsequent login. EX1004, 25:54-56 (“the captured device identity information (ID), including any previously stored Device ID, is compared to identity information that has previously been stored”), 25:7-14 (“the captured device identifying information includes ... [a] first type of device identifying information [that] is a secure, persistent data token that has been previously stored on the user device” and “includes ... data elements that are encrypted, signed, or otherwise secured against modification, and that remain resident on the user device even when it is not accessing a service provider application”). Thus, Varghese discloses that the server identifies information that was previously encrypted or encoded when the server extracts the device identity information in a subsequent authentication attempt.

171. In my opinion, a POSITA would have understood and found it obvious that Varghese discloses the authentication server’s processor is configured to perform this function.

- 9. [1.d.i] and [1.d.ii] - perform a first identification of at least one of the client device and a user of the client device using the first cookie of the first type, wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device**

172. In my opinion, Varghese discloses and renders obvious this limitation.

173. Varghese discloses and renders obvious the authentication server's processor *performs a first identification of ... the client device ... using the first cookie of the first type ... wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device.* Varghese discloses that the secure cookie is used to identify the client device on subsequent log-in attempts. EX1004, 25:7-14 (“the captured device identifying information includes ... [a] first type of device identifying information [that] is a secure, persistent data token that has been previously stored on the user device” and “includes ... data elements that are encrypted, signed, or otherwise secured against modification, and that remain resident on the user device even when it is not accessing a service provider application”), 25:54-56 (“the captured device identity information (ID), including any previously stored Device ID, is compared to identity information that has previously been stored”), 26:4-6 (“a new Device ID token is created for the device ... is sent to the user device and stored thereon ... as a standard cookie”), 25:25-32, 26:9-12 (“[i]f a Device ID was found, it can be updated, e.g., with a new unique bit string, new timestamp, and so forth”), Table 3, 14:67-15:3. The Device ID includes information that can identify the client device, such as, but not limited to “IP addresses, adapter MAC addresses, local time and/or time zone, network

connection speed such as download and/or upload times, microprocessor type and/or processing and/or serial number, and so forth.” EX1004, 25:48-51.

174. Varghese further discloses that the processors *perform a first identification of ... a user of the client device ... using the first cookie of the first type ... wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device.* As shown in Table 3, Varghese discloses that various information is stored in association with the secure cookie, such as post-authentication information, which includes “user identifications.” EX1004, 14:67-15:3 (“Table 3 present a more detailed catalog of device software and hardware characteristics that can be extracted from a device by a browser-hosted process”).

TABLE 3

Example of Request Attributes			
Pre-authentication	Location information	City, State, Country information and confidence factors	
		Connection type	
		Connection speed	
		IP address, routing type, and hop times	
		Internet service provider flag	
		Autonomous system number	
		Carrier name	
		Top-level domain	
		Second-level domain	
		Registering organization	
		A list of anonymizing proxies	
		Hostnames and routers	
		Device information	Secure Cookies
			Flash Cookies
Post-authentication	User information	Digitally signed device	
		Device & display Characteristics:	
		Operating System characteristics	
		Device Characteristics	
		User identifications	
		Valid or not valid user	
		Authentication status	
	Transaction information	Key Value Pairs:	
		Support multiples	
		Keys can be defined using Regular Expressions	
	Values can be defined in ranges		
	Pages accessed		
	Time spent on page		
	Transactions sequences		

EX1004, Table 3 (annotation added). In my opinion, a POSITA would have understood and found it obvious that the encoded information in the secure cookie would be used for the identification of the user because the secure cookie was created as part of the authentication process and Table 3 discloses that after the authentication process, in post-authentication, the user identification is available.

175. In my opinion, a POSITA would have understood and found it obvious that Varghese discloses the authentication server's processor is configured to perform this function.

10. [1.e] - perform a second identification of at least one of the client device and the user of the client device using the second cookie of the second type

176. In my opinion, Varghese discloses and renders obvious this limitation.

177. Varghese discloses and renders obvious the authentication server's processor *perform a second identification of ... the client device ... using the second cookie of the second type*. Varghese discloses that the flash cookie is used to identify the client device on subsequent log-in attempts using a Device ID.

EX1004, 25:7-14 (“the captured device identifying information includes ... [a] first type of device identifying information [that] is a secure, persistent data token that has been previously stored on the user device” and “includes ... data elements that are encrypted, signed, or otherwise secured against modification, and that remain resident on the user device even when it is not accessing a service provider application”), 25:54-56 (“the captured device identity information (ID), including any previously stored Device ID, is compared to identity information that has previously been stored”), 26:4-6 (“a new Device ID token is created for the device ... is sent to the user device and stored thereon ... as a flash cookie”), 26:8-10 (“[i]f a Device ID was found, it can be updated, e.g., with a new unique bit string, new timestamp, and so forth”). The Device ID includes information that can identify the client device, such as, but not limited to “IP addresses, adapter MAC addresses, local time and/or time zone, network connection speed such as

download and/or upload times, microprocessor type and/or processing and/or serial number, and so forth.” EX1004, 25:48-51.

178. Varghese further discloses that the processors *perform a second identification of ... a user of the client device ... using the second cookie of the second type*. As shown in Table 3, Varghese discloses that various information is stored in association with the flash cookie, such as post-authentication information, which includes “user identifications.” EX1004, 14:67-15:3 (“Table 3 present a more detailed catalog of device software and hardware characteristics that can be extracted from a device by a browser-hosted process”).

TABLE 3

Example of Request Attributes

Pre-authentication	Location information	City, State, Country information and confidence factors Connection type Connection speed IP address, routing type, and hop times Internet service provider flag Autonomous system number Carrier name Top-level domain Second-level domain Registering organization A list of anonymizing proxies Hostnames and routers
	Device information	Secure Cookies Flash Cookies Digitally signed device
Post-authentication	User information	Device & display Characteristics: Operating System characteristics Device Characterization User identifications Valid or not valid user Authentication status
	Transaction information	Key Value Pairs: Support multiples Keys can be defined using Regular Expressions Values can be defined in ranges Pages accessed Time spent on page Transactions sequences

EX1004, Table 3 (annotation added). In my opinion, a POSITA would have understood and found it obvious that the encoded information in the flash cookie would be used for the identification of the user because the flash cookie was created as part of the authentication process and Table 3 discloses that after the authentication process, in post-authentication, the user identification is available.

179. In my opinion, a POSITA would have understood and found it obvious that Varghese discloses the authentication server's processor is configured to perform this function.

- 11. [1.f] - perform a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie**

180. In my opinion, Varghese discloses and renders obvious this limitation.

181. Varghese discloses and renders obvious the authentication server's processor configured to *perform a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie*. In my opinion, a POSITA would have understood that a network resource request disclosed in Varghese is an authentication request. §V(C)(4) (limitation 1.b.ii). Varghese discloses that the security server utilizes a security model for determining whether a request presents

a security problem. EX1004, 19:25-28. Table 8 (reproduced and annotated below) illustrates the security model which instructs the security server how to act depending on the presence or absence of either of the secure cookie or the flash cookie. For example, if both the secure cookie and the flash cookie are present, a first determination or score is returned. EX1004, 19:54-56 (“[t]his table returns a score of ‘0’ (a score indicated a low likelihood of fraud) in case all evaluated data items are present and match in connection with a current user request”). In another example, if not all of the data items are present, *e.g.*, the secure cookie is not present but the flash cookie is present, a second determination or score is returned. EX1004, 19:58-65 (“[i]f no data item is present or if all data items do not match, a score of ‘10’ (a score indicated a high likelihood of fraud) is returned ... [i]n case where some data items are present and match while other data items are absent or do not match, this table invokes further checks ... [i]f the retrieved data tokens that were previously stored on a device by this invention, *e.g.*, a secure cookie, a Flash cookie, or Flash data, are not present, a further pattern check is performed.”). If the secure cookie is not present but the flash cookie is present, in my opinion, a POSITA would have understood this to mean that the network resource request associated with the secure cookie was not available but that the network resource request associated with the flash cookie was available.

TABLE 8

Primary device decision table					
Data item					
Secure cookie	Flash cookie	Flash data	Browser characteristics	Operating system characteristics	Score
*	*	*	*	*	0
X	*	*	*	*	PATTERN CHECK
M	*	*	*	*	SECONDARY CHECK
X/M	X	*	*	*	PATTERN CHECK
X/M	M	*	*	*	SECONDARY CHECK
X/M	X/M	X	*	*	PATTERN CHECK
X/M	X/M	M	*	*	SECONDARY CHECK
X/M	X/M	X/M	M	*	SECONDARY CHECK
X/M	X/M	X/M	X/M	M	10

Key:
X = missing;
M = present and mismatched;
* = present and matched

EX1004, Table 8 [annotation added].

Secure cookie is enabled
and flash is disabled
Secure cookie is disabled
and flash is enabled

Only secure cookie came
through successfully.
Only flash cookie came
through successfully.

EX1004, Appendix A (excerpted). Thus, the server makes a pattern check determination based on receipt of a network resource request that includes a flash cookie but not a secure cookie.

182. Varghese also expressly shows that a scenario where a secure cookie (first cookie) is “out of sync” and the flash cookie (second cookie) “is in sync.”

EX1004, Appendix A. In my opinion, a POSITA would have understood this to

mean that a network resource request that includes an out of sync secure cookie is not a valid resource request and that, therefore, the server does not receive a network resource request with a valid or in sync secure cookie. Similarly, in my opinion, a POSITA would have understood that a network resource request that includes an in sync flash cookie is a valid resource request and that, therefore, the server identifies a presence of the flash cookie.

Secure cookie out of
sync and flash is in
sync.
Flash cookie out of
sync and secure cookie
is sync.

Id. (excerpted). Thus, Varghese discloses another determination in which a valid network resource requests for one of the secure cookie and the flash cookie is present and the other is not present.

183. In my opinion, a POSITA would have understood and found it obvious that Varghese discloses the authentication server's processor is configured to perform this function.

12. [1.g] - a memory coupled to the one or more processors and configured to provide the one or more processors with instructions.

184. In my opinion, Varghese discloses and renders obvious this limitation.

185. Varghese discloses *a memory coupled to the one or more processors and configured to provide the one or more processors with instructions*. Varghese's systems server 1306 "includes a CPU, RAM memory, disc or other database memory 1308, communication interfaces, optional user interface equipment, and the like." EX1004, 8:51-53. Further, in my opinion, a POSITA would have understood that system server's 1306 memory would need to include instructions that are provided to the authentication server's processor otherwise the processor would not know how to perform the generation and validation steps described above.

D. Claim 3 - The system of claim 1 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.

186. In my opinion, Varghese discloses and renders this claim obvious.

187. As I explain above in Section V(C), Varghese discloses and renders obvious *the system of claim 1*. See §V(C) (claim 1).

188. The client device is identifiable by a first cookie in a first client device browser storage area. §V(C)(7), (9) (limitations 1.b.v, 1.d.i), The client device is also identifiable by a second cookie in a second client device browser storage area. §V(C)(7), (11) (limitations 1.b.v, 1.e).

E. Claim 4 - The system of claim 1 wherein the determination comprises detection of pharming.

189. In my opinion, Varghese discloses and renders this claim obvious.

190. As I explain above in Section V(C), Varghese discloses and renders obvious *the system of claim 1*. See §V(C) (claim 1).

191. A POSITA would have been familiar with the problem of domain spoofing, or pharming ((EX1015)“The Pharming Guide (part 2)” Dec 14, 2004 available at: <http://www.technicalinfo.net/papers/Pharming2.html>). And, in my opinion, a POSITA would have been motivated to add detection of pharming to Varghese. A POSITA would have been motivated by server load balancing to implement the Flash cookie at the client device to include a server identifier, such as an IP address, for the authentication server in order to prevent possible domain spoofing. *See, e.g.*, EX1013. In this manner, when the user attempts to use the Flash cookie to authenticate themselves at the authentication server, the authentication server verifies whether the cookie includes both the URL of the authentication server and the IP address of the authentication server, as an additional check to determine whether a malicious actor spoofed its URL, *i.e.*, perform a detection of pharming, based on the Flash cookie. If the authentication server determines that the URL contained in the cookie is correct but the IP address contained in the cookie is incorrect, then it has detected pharming.

F. Claim 5 - The system of claim 1 wherein in response to the performed determination, the one or more processors are configured to cause a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.

192. In my opinion, Varghese discloses and renders this claim obvious.

193. As I explain above in Section V(C), Varghese discloses and renders obvious *the system of claim 1*. See §V(C) (claim 1).

194. In the case where the determination is performed responsive to a subsequent login request, Varghese discloses that one of the secure token or the flash token is replaced by the authentication server upon each login. EX1004, 26:13-14 (“[a] feature of the invention relates to the replacement of the cookie on the user’s machine upon each login.”). Thus, the replacement token is a third cookie of one of the first or second type.

G. Claim 6

195. Varghese discloses and renders claim 6 obvious. See §V(C) (claim 1).

1. [6.pre] A method, comprising

196. Varghese discloses and renders obvious the preamble. Above in Section §V(C), I show that Varghese discloses a system that performs a method. See §V(C) (claim 1).

2. **[6.a.i] receiving a network resource request from a client device, wherein the network resource request corresponds to a first cookie of a first type that was caused to be stored to the client device during a first previous network session,**

197. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(3) (limitation 1.b.i).

3. **[6.a.ii] wherein the first cookie of the first type was caused to be stored to the client device at least in part by causing the client device to initiate a set of network resource requests determined during the first previous network session,**

198. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(4) (limitation 1.b.ii).

4. **[6.a.iii] wherein the client device initiating the set of network resource requests caused data representative of the set of network resource requests to be stored at the client device,**

199. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(5) (limitation 1.b.iii).

5. **[6.a.iv] wherein a second cookie of a second type different from the first type was caused to be stored at the client device during a second previous network session, and**

200. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(6) (limitation 1.b.iv).

6. **[6.a.v] wherein the first cookie of the first type is stored in a first client device browser storage area and the second cookie of the second type is stored in a second client device browser storage area different from the first client device browser storage area;**

201. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(7) (limitation 1.b.v).

7. **[6.b] based at least in part on the network resource request from the client device corresponding to the first cookie of the first type caused to be stored at the client device during the first previous network session, determining information that was encoded and stored in the client device; and**

202. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(8) (limitation 1.c).

8. **[6.c.i] and [6.c.ii] performing a first identification of at least one of the client device and a user of the client device using the first cookie of the first type, wherein the first identification is performed using the first cookie of the first type at least in part by using the determined information that was encoded and stored in the client device**

203. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(9) (limitation 1.d.i and 1.d.ii).

9. **[6.d] performing a second identification of at least one of the client device and the user of the client device using the second cookie of the second type**

204. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(10) (limitation 1.e).

- 10. [6.e] performing a determination based at least in part on (1) a presence of a network resource request associated with one of the first cookie and the second cookie, and (2) an absence of a network resource request associated with the other of the first cookie and the second cookie.**

205. In my opinion, Varghese discloses and renders obvious this limitation.

See §V(C)(11) (limitation 1.f).

- H. Claim 8 - The method of claim 6 wherein the client device is identifiable based at least in part on a presence of either one of the first cookie and the second cookie in the respective first and second client device browser storage areas.**

206. Varghese discloses and renders obvious this claim. *See* §§V(D), (G)

(claims 3, 6).

- I. Claim 9 - The method of claim 6 wherein the determination comprises detection of pharming.**

207. Varghese discloses and renders obvious this claim. *See* §§V(E), (G)

(claims 4, 6).

- J. Claim 10 - The method of claim 6 wherein in response to the performed determination, further comprising causing a third cookie to be stored at the client device, wherein the third cookie is one of the first type and the second type.**

208. Varghese discloses and renders obvious this claim. *See* §§V(F)-(G)

(claims 5-6).

209. To the extent it is asserted that the portions of Varghese that I have cited above relate to different, incompatible embodiments (which they do not), in

my opinion, it would have been obvious to a POSITA to combine such embodiments into a single system for client-server communications at least because such embodiments are described in the same prior art reference, are fully compatible with each other, and could be combined with minimal effort to achieve predictable results.

210. To the extent it is asserted that Varghese does not disclose any limitation of the challenged claims, in my opinion, such limitation would have nonetheless been obvious to a POSITA in light of Varghese and a POSITA's knowledge. Practicing any limitation of the challenged claims in light of Varghese would have been within the knowledge and skill of a POSITA, would have required minimal effort, would have yielded predictable results, would have been fully compatible with the Varghese system, and would have been a mere design choice. Motivation to do so arises from at least common sense and the disclosures of Varghese that I have set forth above.

VI. GROUND 3: VARGHESE IN VIEW OF HINTON RENDERS OBVIOUS CLAIMS 2 AND 7

A. Overview of Varghese

211. My explanation of the Varghese reference is above and herein incorporated by reference into Ground 3. *See* §V(A).

B. Overview of Hinton

212. My explanation of the Hinton reference is above and herein incorporated by reference into Ground 3. *See* §IV(A).

C. Motivation to Combine

213. In my opinion, a POSITA would have been motivated to modify the teachings of Varghese with Hinton because Varghese teaches an authentication system for providing protection against identity theft over a computer network and Hinton teaches a means for authenticating users across systems situated in different computer domains. *See* §§IV(A), V(A).

214. In my opinion, a POSITA would have been motivated to modify Varghese's computing environment such that Varghese's systems (*e.g.*, systems 1302, 1304, 1306) can be distributed across different domains. Such a combination involves a combination of prior art elements (*e.g.*, combining Hinton's distributed system with Varghese's authentication system) according to known methods to yield predictable results (*e.g.*, such as in the situation where Varghese's authentication process is utilized in a business-to-business environment). EX1005, ¶10. In my opinion, a POSITA would have been motivated by the teachings of Hinton to modify Varghese such that system 1302 is in a different domain than system 1306, such as, for example, in a business-to-business use case. Moreover, in the timeframe, a POSITA would have known that a single organization was known

to use multiple domains with multiple servers for providing content. EX1012.

Thus, it was well-known that a business would need to support multiple domains, so a POSITA would have been motivated to modify Varghese to do so (as taught by Hinton). Therefore, it was obvious to a POSITA that the provision of cookies in a single domain would be used by another domain.

215. In my opinion, a POSITA would have had a reasonable expectation of success in making the proposed combination given the teachings of Varghese and Hinton. Such a combination would have required minimal modifications to Varghese, as the use of cookies to identify users when interacting with servers (regardless of the domain) is well known in the art and would have been well within the skillset of a POSITA. The proposed combination would not have required undue experimentation and would have yielded the predictable result.

D. Claim 2 - The system of claim 1 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.

216. In my opinion, Varghese in view of Hinton discloses and renders this claim obvious. *See* §V(C) (claim 1), §IV(D) (claim 2).

217. Varghese discloses an authentication system in which systems 1302, 1304, and 1306 receive network resource requests from a client device. §V(C)(3) (limitation 1.b.i); EX1004, Fig. 13A, (“[t]he present invention enables application service providers score risk for each online login and transaction and to increase

authentication security in real time”). As part of the process, server 1306, for example, creates cookies—secure cookie, flash cookie—that can be used to authenticate individuals. §§V(C)(4), (6) (limitations 1.b.ii, 1.b.iv). Server 1302 may utilize the cookies in authenticating the user (“system 1302 does not itself perform pre-authentication processing, but does performs[sic] ‘post-authentication services’”). Although Varghese does not explicitly disclose that server 1302 and server 1306 are part of the same domain, Hinton discloses a computing environment in which cookies can be used to authenticate users at servers located in different domains. §IV(D) (claim 2).

218. In my opinion, a POSITA would have been motivated by the teachings of Hinton to modify Varghese such that system 1302 is in a different domain than system 1306 and had a reasonable expectation of success. §V(B) (motivation to combine). Accordingly, Varghese in view of Hinton yields a computing environment in which the user is provided with cookies generated by a first system located in a first domain (*e.g.*, system 1306) and can use those cookies to verify themselves at a second system in a second domain (*e.g.*, system 1304).

E. Claim 7 - The method of claim 6 wherein the client device is identified by a server in a domain different from a server that caused the two different types of cookies to be stored to the client device during the first and second previous network sessions.

219. In my opinion, Varghese in view of Hinton discloses and renders this claim obvious. *See* §V(G) (claim 6), §VI(D) (claim 2).

VII. SECONDARY CONSIDERATIONS

220. I am not aware of any relevant information related to secondary considerations of obviousness. For example, I am not aware of any attempt to commercialize the claimed technology in a product or through licensing. To the extent Patent Owner comes forward with any such information, I will consider it.

VIII. AVAILABILITY FOR CROSS-EXAMINATION

221. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during a time as mutually agreed between the parties.

IX. RIGHT TO SUPPLEMENT

222. I reserve the right to supplement my opinions in the future to respond to any arguments that the Patent Owner raises. This declaration represents only those opinions that I have formed to date. I reserve the right to revise, supplement, and/or amend my opinions stated herein based on new information that becomes available to me and on my continuing analysis of the materials already provided.

X. CONCLUSION

223. I may utilize the documents cited and/or listed herein, or portions of those documents, as exhibits at any hearing or trial in this proceeding. I may further prepare and use exhibits that summarize portions of my testimony or key terms or concepts presented therein, or other demonstrative exhibits, at any hearing or trial in this proceeding.

224. I reserve the right to supplement my testimony and this report in response to any judicial determinations, in response to the arguments expressed by the Patent Owner or the opinions of the Patent Owner's experts in this proceeding, and/or in light of additional evidence or testimony brought forth at trial or otherwise brought to my attention after the date of my signature on the cover page.

APPENDIX 1

Craig E. Wills

**Professor
Computer Science Department
Worcester Polytechnic Institute
Worcester, MA 01609
cew@wpi.edu**

Education

Ph.D. Computer Science. Purdue University. May, 1988. Dissertation title “Service Execution in a Distributed Environment.” Committee: J.T. Korb (advisor), D. Comer, P. Mehrotra and R. Stansifer

M.S. Computer Science. Purdue University. May, 1984

B.S. Computer Science. University of Nebraska. May, 1982

Professional Experience

Computer Science Department, Worcester Polytechnic Institute. Professor. July, 2022 – present

Department of Computer Science and Engineering, University of Bologna, Italy. Visiting Professor. Host: Michele Colajanni. March, 2023 – May, 2023

Computer Science Department, Worcester Polytechnic Institute. Professor and Department Head. January, 2011 – June, 2022

Computer Science Department, Worcester Polytechnic Institute. Professor. July, 2010 – December, 2010

Computer Science Department, Worcester Polytechnic Institute. Associate Professor. July, 1996 – June, 2010

Cisco Systems, Inc. Boxborough, Massachusetts. Visiting Faculty, Network Management Technology Group. August, 2004 – July, 2005

School of Mathematical and Computing Sciences, Victoria University of Wellington, New Zealand. Visiting Professor. July, 1997 – June, 1998

Computer Science Department, Worcester Polytechnic Institute. Assistant Professor. August, 1990 – June, 1996

AT&T Bell Laboratories, Middletown, New Jersey. Design and development of an automation tool for network management applications. May, 1988 – August, 1990

Department of Computer Science, Purdue University. Instructor, programmer, research assistant and grader. August, 1982 – May, 1988

Research Interests

Internet application performance and measurement; security/privacy; distributed systems; networking; operating systems; human-computer interaction; Computer Science education and workforce; data-driven analysis of higher education, geography and sports.

Publications

Citations 6210 total citations with h-index of 34 based upon information available from Google Scholar Profile. March 2024

Books and Book Chapters

1. Craig E. Wills. Process synchronization and IPC. In Teo Gonzalez and Jorge L. Diaz-Herrera, editors, *Computer Science and Software Engineering Computing Handbook*. CRC Press, 2013
2. Craig E. Wills. Process synchronization and IPC. In Allen Tucker, editor, *Computer Science Handbook, Second Edition*, chapter 84, pages 84–1–84–22. CRC Press, 2004
3. Craig E. Wills, Kirstin Cadwell, and William Marrs. Customization in a UNIX computing environment. In Eric Anderson, Mark Burgess, and Alva Couch, editors, *Selected Papers in Network and System Administration*, pages 203–209. John Wiley & Sons Ltd, 2002. Compilation of significant contributions to the field of system administration
4. Craig E. Wills. Process synchronization and IPC. In Allen Tucker, editor, *Handbook of Computer Science and Engineering*, chapter 79, pages 1725–1746. CRC Press, 1996
5. Craig E. Wills. A model for executing computations in a distributed environment. In T.L. Casavant and M. Singhal, editors, *Readings in Distributed Computing Systems*, pages 116–132. IEEE Computer Society Press, 1994

Refereed Journals

1. Craig E. Wills. The competitiveness of games in professional sports leagues. *Journal of Sports Analytics*, 3(2):103–117, July 2017.
2. Craig E. Wills and Mihajlo Zeljkovic. A personalized approach to web privacy—awareness, attitudes and actions. *Information Management and Computer Security*, 19(1):53–73, 2011.
3. Balachander Krishnamurthy and Craig E. Wills. On the leakage of personally identifiable information via online social networks. *ACM SIGCOMM Computer Communications Review*, 40(1):112–117, January 2010. Accepted for publication based on selection as one of the two best papers at the Workshop on Online Social Networks held in conjunction with ACM SIGCOMM 2009 Conference. This is a minor revision of the original workshop publication.
4. Hao Shang and Craig E. Wills. Making better use of all those TCP ACK packets. *ISAST Transactions on Communications and Networking*, 1(1):32–41, 2007.

5. Hao Shang and Craig E. Wills. Piggybacking related names to improve DNS performance. *Computer Networks*, 50:1733–1748, August 2006.
6. Chi-Hung Chi, Maarten van Steen, and Craig E. Wills, editors. *Web Content Caching and Distribution: 9th International Workshop, WCW 2004, Beijing, China, October 18-20, 2004. Proceedings*, volume 3293 of *Lecture Notes in Computer Science*. Springer, 2004
7. Craig E. Wills, Gregory Trott, and Mikhail Mikhailov. Using bundles for web content delivery. *Computer Networks*, 42(6):797–817, August 2003.
8. David Finkel, Craig E. Wills, Michael Ciaraldi, Kevin Amorin, Adam Covati, and Michael Lee. An applet-based anonymous distributed computing system. *Internet Research: Electronic Networking Applications and Policy*, 11(1):35–41, 2001.
9. Craig E. Wills and Mikhail Mikhailov. Studying the impact of more complete server information on web caching. *Computer Communications*, 24(2):184–190, February 2001. Published as the Proceedings of the 5th International Web Caching and Content Delivery Workshop.
10. Balachander Krishnamurthy and Craig E. Wills. Analyzing factors that influence end-to-end web performance. *Computer Networks*, 33(1-6):17–32, June 2000. Published as the Proceedings of the Ninth International World Wide Conference.
11. Craig E. Wills, Dorothy Deremer, Renee A. McCauley, and Linda Null. Studying the use of peer learning in the introductory computer science curriculum. *Computer Science Education*, 9(2):71–88, August 1999.
12. Craig E. Wills and Mikhail Mikhailov. Towards a better understanding of web resources and server responses for improved caching. *Computer Networks*, 31(11-16):1231–1243, May 1999. Published as the Proceedings of the Eighth International World Wide Conference.
13. David Finkel, Craig E. Wills, Brian Brennan, and Chris Brennan. Dtriblets: Java-based distributed computing on the web. *Internet Research: Electronic Networking Applications and Policy*, 9(1):35–40, 1999. Paper awarded Citation of Excellence (made to less than 10% of reviewed papers) by ANBAR Electronic Intelligence.
14. C.E. Wills, D.C. Brown, B. Dunskus, and J. Kemble. Evaluating network serviceability. *Computer Networks and ISDN Systems*, 30(24):2283–2291, December 1998.
15. Balachander Krishnamurthy and Craig E. Wills. Piggyback server invalidation for proxy cache coherency. *Computer Networks and ISDN Systems*, 30(1-7):185–193, April 1998. Published as the Proceedings of the Seventh International World Wide Conference.
16. Marton E. Balazs, David C. Brown, Peter Bastien, and Craig E. Wills. Graphical presentation of designs: A knowledge intensive design approach. In M. Mantyla, S. Finger, and T. Tomiyama, editors, *Knowledge Intensive CAD*, volume II, pages 173–188. Chapman & Hall, 1997.

17. Craig E. Wills. Process synchronization and IPC. *ACM Computing Surveys*, 28(1):209–211, March 1996. 50th-anniversary issue
18. Craig E. Wills and David Finkel. Scalable approaches to load sharing in the presence of multicasting. *Computer Communications*, 18(9):620–630, September 1995.
19. Craig E. Wills and David Finkel. Experience with peer learning in an introductory computer science course. *Computer Science Education*, 5(2):165–187, 1994.

Professional Publications

1. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2024. *Computing Research News*, 36(1), January 2024.
2. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2023. *Computing Research News*, 35(1), January 2023.
3. Craig E. Wills. 2022 computer science tenure-track faculty hiring outcomes. *Computing Research News*, 34(10), November 2022.
4. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2022. *Computing Research News*, 34(1), January 2022.
5. Craig E. Wills. Updated analysis of current and future computer science needs via advertised faculty searches for 2021. *Computing Research News*, 33(1), January 2021.
6. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2021. *CRA Bulletin*, December 2, 2020.
7. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2020. *Computing Research News*, 32(1), January 2020.
8. Craig E. Wills. 2019 computer science tenure-track faculty hiring outcomes. *Computing Research News*, 31(10), November 2019.
9. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2019. *Computing Research News*, 31(1), January 2019.
10. Craig E. Wills. Outcomes of advertised computer science faculty searches for 2018. *Computing Research News*, 30(7), August 2018.
11. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2018. *Computing Research News*, 30(1), January 2018.
12. Craig E. Wills. Outcomes of advertised computer science faculty searches for 2017. *Computing Research News*, 29(10), November 2017.
13. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2017. *Computing Research News*, 29(1), January 2017.

14. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2016. *Computing Research News*, 28(1), January 2016.
15. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches. *Computing Research News*, 27(1), January 2015.
16. David Finkel and Craig E. Wills. Peer learning assistants in an introductory computer science course. *INPUT, A Newsletter for Computer Science Educators*, pages 5–6, Spring 1995

Refereed Conference Publications

1. Craig E. Wills and Doruk C. Uzunoglu. What ad blockers are (and are not) doing. In *Proceedings of the IEEE Workshop on Hot Topics in Web Systems and Technologies*, Washington, DC USA, October 2016.
2. Craig E. Wills and Can Tatar. Understanding what they do with what they know. In *Proceedings of the Workshop on Privacy in the Electronic Society*, Raleigh, NC USA, October 2012. Acceptance rate: 30%. See technical report for a more detailed version of the paper with illustrative examples.
3. Murad Kaplan, Mihajlo Zeljkovic, Mark Claypool, and Craig E. Wills. How's my network? predicting performance from within a web browser sandbox. In *Proceedings of the IEEE Conference on Local Computer Networks*, pages 525–532, Clearwater, FL USA, October 2012. Acceptance rate: 29%.
4. Wei Zhang and Craig E. Wills. Consideration of processing costs in placing clients of web-based services. In *Proceedings of the IEEE GLOBECOM Conference*, Houston, TX USA, December 2011. Acceptance rate: 37%.
5. Balachander Krishnamurthy, Konstantin Naryshkin, and Craig E. Wills. Privacy leakage vs. protection measures: The growing disconnect. In *Proceedings of the Web 2.0 Security and Privacy Workshop*, pages 1–10, Oakland, CA USA, May 2011. Acceptance rate: 27% for full papers.
6. Balachander Krishnamurthy and Craig E. Wills. Privacy leakage in mobile online social networks. In *Proceedings of the Workshop on Online Social Networks*, pages 1–9, Boston, MA USA, June 2010. USENIX. Acceptance rate: 33%.
7. Wei Zhang, Hangwei Qian, Craig E. Wills, and Michael Rabinovich. Agile resource management in a virtualized data center. In *Proceedings of the First Joint WOSP/SIPEW International Conference on Performance Engineering*, San Jose, California USA, January 2010. ACM. Acceptance rate: <25%.
8. Artur Janc, Craig E. Wills, and Mark Claypool. Network performance evaluation in a web browser. In *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems*, pages 370–377, Cambridge, MA USA, November 2009.

9. Alan Ritacco, Craig E. Wills, and Mark Claypool. How's my network? a Java approach to home network measurement. In *Proceedings of the IEEE International Conference on Computer Communications and Networks*, pages 1–7, San Francisco, CA USA, August 2009. Acceptance rate: 30%.
10. Balachander Krishnamurthy and Craig E. Wills. On the leakage of personally identifiable information via online social networks. In *Proceedings of the Workshop on Online Social Networks in conjunction with ACM SIGCOMM Conference*, pages 7–12, Barcelona, Spain, August 2009. Acceptance rate: 37%.
11. Balachander Krishnamurthy and Craig E. Wills. Privacy diffusion on the web: A longitudinal perspective. In *Proceedings of the World Wide Web Conference*, pages 541–550, Madrid, Spain, April 2009. Acceptance rate: 13%.
12. Balachander Krishnamurthy and Craig E. Wills. Characterizing privacy in online social networks. In *Proceedings of the Workshop on Online Social Networks in conjunction with ACM SIGCOMM Conference*, pages 37–42, Seattle, WA USA, August 2008. Acceptance rate: 35%.
13. Hangwei Qian, Elliot Miller, Wei Zhang, Michael Rabinovich, and Craig E. Wills. Agility in virtualized utility computing. In *Proceedings of the 3rd International Workshop on Virtualization Technology in Distributed Computing*, pages 1–8, Reno, NV USA, November 2007. Held in conjunction with ACM/IEEE Super Computing Conference.
14. Mark Claypool, Robert Kinicki, and Craig Wills. Treatment-based traffic signatures. In *Proceedings of the IMRG (IETF Internet Measurement Research Group) Workshop on Application Classification and Identification (WACI)*, Cambridge, MA USA, October 2007.
15. Balachander Krishnamurthy, Delfina Malandrino, and Craig E. Wills. Measuring privacy loss and the impact of privacy protection in web browsing. In *Proceedings of the Symposium on Usable Privacy and Security*, pages 52–63, Pittsburgh, PA USA, July 2007. Acceptance rate: 32%.
16. Paul J. Timmins, Sean McCormick, Emmanuel Agu, and Craig E. Wills. Characteristics of mobile Web content. In *Proceedings of the First IEEE Workshop on Hot Topics in Web Systems and Technologies*, pages 1–10, Boston, MA USA, November 2006.
17. Balachander Krishnamurthy and Craig E. Wills. Generating a privacy footprint on the Internet. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, pages 65–70, Rio de Janeiro, Brazil, October 2006. Acceptance rate: 22%.
18. Balachander Krishnamurthy and Craig Wills. Cat and mouse: Content delivery tradeoffs in web access. In *Proceedings of the International World Wide Web Conference*, pages 337–346, Edinburgh, Scotland, May 2006. Acceptance rate: 11%.
19. Craig E. Wills, Mikhail Mikhailov, and Hao Shang. Inferring relative popularity of Internet applications by actively querying DNS caches. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, pages 78–90, Miami, Florida, October 2003. Acceptance rate: 30%.

20. Mikhail Mikhailov and Craig E. Wills. Evaluating a new approach to strong web cache consistency with snapshots of collected content. In *Proceedings of the Twelfth International World Wide Web Conference*, pages 599–608, Budapest, Hungary, May 2003. Acceptance rate: 13%.
21. Balachander Krishnamurthy, Craig Wills, Yin Zhang, and Kashi Vishwanath. Design, implementation, and evaluation of a client characterization driven web server. In *Proceedings of the Twelfth International World Wide Web Conference*, pages 138–147, Budapest, Hungary, May 2003. Acceptance rate: 13%.
22. Balachander Krishnamurthy, Craig Wills, and Yin Zhang. Preliminary measurements on the effect of server adaptation for web content delivery. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop*, pages 323–324, Marseille, France, November 2002. Acceptance rate: 45%. Short paper version accepted.
23. Mikhail Mikhailov and Craig Wills. Exploiting object relationships for deterministic web object management. In *Proceedings of the 7th International Workshop on Web Content Caching and Distribution*, pages 1–16, Boulder, Colorado, August 2002. Acceptance rate: 29%.
24. Balachander Krishnamurthy and Craig Wills. Improving web performance by client characterization driven server adaptation. In *Proceedings of the Eleventh International World Wide Web Conference*, pages 305–316, Honolulu, Hawaii, May 2002. Acceptance rate: 16%.
25. Balachander Krishnamurthy, Craig Wills, and Yin Zhang. On the use and performance of content distribution networks. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop*, pages 169–182, San Francisco, November 2001.
26. Mark Claypool, David Finkel, and Craig E. Wills. An open source laboratory for operating systems projects. In *ACM SIGCSE/SIGCUE Conference on Innovation and Technology in Computer Science Education*, pages 145–148, Canterbury, England, June 2001. Acceptance rate: 30%.
27. Craig E. Wills, Mikhail Mikhailov, and Hao Shang. N for the price of 1: Bundling web objects for more efficient content delivery. In *Proceedings of the Tenth International World Wide Web Conference*, pages 257–265, Hong Kong, May 2001. Acceptance rate: 20%.
28. Craig E. Wills and Mikhail Mikhailov. Studying the impact of more complete server information on web caching. In *Proceedings of the 5th International Web Caching and Content Delivery Workshop*, pages 184–190, Lisbon, Portugal, May 2000. Acceptance rate: 46%.
29. David Finkel, Craig E. Wills, Kevin Amorin, Adam Covati, and Michael Lee. An applet-based approach to large-scale distributed computing. In *Proceedings of the International Network Conference*, pages 175–182, Plymouth, United Kingdom, July 2000. Acceptance rate: 60%.
30. Michael J. Ciaraldi, David Finkel, and Craig E. Wills. Risks in anonymous distributed computing systems. In *Proceedings of the International Network Conference*, pages 193–200, Plymouth, United Kingdom, July 2000. Acceptance rate: 60%.

31. Balachander Krishnamurthy and Craig E. Wills. Analyzing factors that influence end-to-end web performance. In *Proceedings of the Ninth International World Wide Web Conference*, pages 17–32, Amsterdam, Netherlands, May 2000. Acceptance rate: 20%.
32. David C. Brown, Isabel Cruz, David Finkel, Robert E. Kinicki, and Craig E. Wills. Experiences with the webware, interfaces and networking experimental laboratory. In *Proceedings of the ACM SIGCSE Conference*, pages 387–391, Austin, TX, March 2000. Acceptance rate: 35%.
33. James F. Carlson, David V. Esposito, Nathaniel J. Springer, David Finkel, and Craig E. Wills. Applet-based distributed computing on the web. In *Proceedings of the Workshop on Distributed Computing on the Web*, Rostock, Germany, June 1999
34. Balachander Krishnamurthy and Craig E. Wills. Proxy cache coherency and replacement—towards a more complete picture. In *Proceedings of the 19th IEEE International Conference on Distributed Computing Systems*, pages 332–339, Austin, TX, June 1999. Acceptance rate: 33%.
35. Craig E. Wills and Mikhail Mikhailov. Towards a better understanding of web resources and server responses for improved caching. In *Proceedings of the Eighth International World Wide Web Conference*, pages 153–165, Toronto, Canada, May 1999. Acceptance rate: 16%.
36. Craig E. Wills and Mikhail Mikhailov. Examining the cacheability of user-requested web resources. In *Proceedings of the 4th International Web Caching Workshop*, pages 78–87, San Diego, CA, March/April 1999. Acceptance rate: 51%.
37. Craig E. Wills and Paul Thomas. Exploiting a network charging model to reduce web costs. In *Proceedings of the AusWeb99—The Fifth Australian World Wide Web Conference*, Ballina, NSW Australia, April 1999. Acceptance rate: 50%. Full paper accepted.
38. John H. Hine, Craig E. Wills, Anja Martel, and Joel Sommers. Combining client knowledge and resource dependencies for improved world wide web performance. In *Proceedings of the INET '98 Conference*, Geneva, Switzerland, July 1998. Internet Society.
39. Brian Brennan, Chris Brennan, David Finkel, and Craig E. Wills. Java-based load distribution on the world wide web. In *Proceedings of the International Network Conference*, pages 9–14, Plymouth, United Kingdom, July 1998.
40. Balachander Krishnamurthy and Craig E. Wills. Piggyback server invalidation for proxy cache coherency. In *Proceedings of the Seventh International World Wide Web Conference*, pages 185–193, Brisbane, Australia, April 1998.
41. Craig E. Wills. Group-based software engineering in an introductory computer science course. In *Proceedings of the International Conference on Software Engineering: Education & Practice Conference*, pages 26–33, Dunedin, New Zealand, January 1998. IEEE Computer Society Press.

42. Balachander Krishnamurthy and Craig E. Wills. Study of piggyback cache validation for proxy caches in the world wide web. In *Proceedings of the Symposium on Internet Technologies and Systems*, pages 1–12. USENIX Association, December 1997. Acceptance rate: 27%.
43. Craig E. Wills and David Finkel. Study of a group project model in computer science. In *Proceedings of the ASEE/IEEE Frontiers in Education Conference*, pages 299–303, Pittsburgh, PA, November 1997.
44. I. Russell, M. Dickerson, G. Scragg, M. Towhidnejad, and C. Wills. Novel approaches to the introductory computer science courses. In *Proceedings of the Second Annual Consortium for Computing in Small Colleges: Northeastern Conference*, pages 170–175, Boston, MA, April 1997
45. C.E. Wills, D. Cordes, D. Deremer, B.J. Klein, R.A. McCauley, and L. Null. Application of peer learning to the introductory computer science curriculum. In *Proceedings of the ACM SIGCSE Conference*, pages 373–374, San Jose, CA, March 1997. Panel presentation
46. C.E. Wills and P.F. Bastien. The influence of resource dependencies on distributed scheduling policies for load sharing. In *Proceedings of the International Conference on Parallel and Distributed Systems*, pages 104–109, Dijon, France, September 1996.
47. Craig E. Wills, Robert E. Kinicki, and David Finkel. Networking projects in the undergraduate curriculum. *Journal of Computing in Small Colleges*, 11(4):238–245, March 1996. Based on a presentation at the First Annual Northeastern Small College Computing Conference. West Hartford, CT. April 1996
48. M.E. Balazs, D.C. Brown, P. Bastien, and C.E. Wills. How to present designs. In *Proceedings of the Second Workshop Knowledge Intensive CAD*, Pittsburgh, PA, September 1996. IFIP Working Group 5.2
49. David Finkel and Craig E. Wills. Computer supported peer learning in an introductory computer science course. In *ACM SIGCSE/SIGCUE Conference on Integrating Technology into Computer Science Education*, pages 55–56, Barcelona, Spain, June 1996
50. D.C. Brown, C.E. Wills, B. Dunskus, and J. Kemble. Tennis: A computer network ease of service evaluation system. In *Proceedings of the International Joint Conference on Artificial Intelligence Workshop on AI in Distributed Information Networks*, Montreal, Canada, August 1995
51. Craig E. Wills, Gregory J. Snyder, and Christopher Kmiec. Persistent information retrieval on the Internet. In *Proceedings of the IASTED/ISMM International Conference on Intelligent Information Management Systems*, pages 152–155, Washington, D.C., June 1995
52. Craig E. Wills, D. Giampaolo, and M. Mackovitch. Experience with an interactive attribute-based user information environment. In *Proceedings of the Fourteenth Annual IEEE International Phoenix Conference on Computers and Communications*, pages 359–365, Phoenix, AZ, March 1995

53. Craig E. Wills and Surendar Chandra. Adaptive resource management. In *Proceedings of The International Workshop on Modeling, Analysis and Simulation of Computers and Telecommunications Systems (MASCOTS'95)*, pages 173–177, Durham, NC, January 1995. Acceptance rate: 45%
54. Craig E. Wills, David Finkel, Michael A. Gennert, and Matthew O. Ward. Peer learning in an introductory computer science course. In *Proceedings of the ACM SIGCSE Conference*, pages 309–313, Phoenix, AZ, March 1994
55. Craig E. Wills, Joachim Heck, and Ramin Taraz. Visualization of a user's information space. In *Proceedings of the Computer Science Conference*, pages 94–101, Phoenix, AZ, March 1994. ACM. Acceptance rate: 40%
56. Craig E. Wills, Kirstin Cadwell, and William Marrs. Customization in a unix computing environment. In *Proceedings of the 7th USENIX System Administration Conference*, pages 43–49, Monterey, CA, November 1993
57. J. CaraDonna, N. Paciorek, and C.E. Wills. Measuring lock performance in multiprocessor operating system kernels. In *Proceedings of the Fourth USENIX Symposium on Experiences with Distributed and Multiprocessor Systems*, pages 43–49, San Diego, CA, September 1993
58. Craig E. Wills, Kirstin Cadwell, and William Marrs. Sharing customization in a campus computing environment. In *HCI International '93*, pages 105–115, Orlando, FL, August 1993
59. Craig E. Wills and David Finkel. Load sharing using multicasting. In *Proceedings of the Twelfth Annual IEEE International Phoenix Conference on Computers and Communications*, pages 303–309, Phoenix, AZ, March 1993
60. Craig E. Wills and Shanti Suresh. Resource-driven resource location. In *Proceedings of the 26th Hawaii International Conference on System Sciences*, pages 80–89, Maui, Hawaii, January 1993
61. Craig E. Wills. Strategies for using multicasting to locate resources. In *Proceedings IEEE 16th Conference on Local Computer Networks*, pages 589–598, Minneapolis, MN, October 1991
62. Craig E. Wills. A service execution mechanism for a distributed environment. In *Proceedings of the 9th IEEE International Conference on Distributed Computing Systems*, pages 326–334, Newport Beach, CA, June 1989. Acceptance rate: 33%
63. Craig E. Wills. Locating distributed information. In *Proceedings IEEE Infocom '89*, pages 303–311, Ottawa, Canada, April 1989
64. John T. Korb and Craig E. Wills. Command execution in a heterogeneous environment. In *SIGCOMM '86 Symposium*, pages 68–74, Stowe, VT, August 1986. ACM

Invited Conference Publications

1. Craig Wills, Mark Claypool, Artur Janc, and Alan Ritacco. Development of a user-centered network measurement platform. In *Proceedings of the ISMA Workshop on Active Internet Measurements*, La Jolla, CA USA, February 2010. Invited participant. Sponsored by CAIDA
2. Mark Claypool, Robert Kinicki, and Craig Wills. Treatment-based traffic signatures. In *Proceedings of the IMRG (IETF Internet Measurement Research Group) Workshop on Application Classification and Identification (WACI)*, Cambridge, MA USA, October 2007.
3. Mark Claypool, Robert Kinicki, and Craig Wills. Research resources for network application studies. In *Proceedings of the National Science Foundation Computing Research Infrastructure 2007 PI Meeting*, pages 143–147, Boston, MA USA, June 2007. Also available as Technical Report WPI-CS-TR-07-03.
4. Craig E. Wills and Mikhail Mikhailov. Characterizing web resources and server responses to better understand the potential of caching. In *Web Characterization Workshop*, Cambridge, MA, November 1998. World Wide Web Consortium.
5. Balachander Krishnamurthy and Craig E. Wills. Piggyback cache validation for proxy caches in the world wide web. In *Proceedings of the 2nd Web Caching Workshop*, Boulder, CO, June 1997. National Laboratory for Applied Network Research.
6. Craig E. Wills. User interface design for the engineer. In *Proceedings of Electro/94 International*, pages 415–419, Boston, MA, May 1994

Other Accepted Conference Publications

1. Craig E. Wills. User and resource efficient access to information in mobile and web domains. In *Proceedings of the Second Annual Conference on Telecommunications R&D in Massachusetts*, Lowell, MA, March 1996
2. C. Council, E. Felton, C. Johnson, R. Mason, R. Rubinstein, and C. Wills. A virtual reality world builder. In *Proceedings of CONVERGENCE: The Fifth Biennial Symposium on the Arts and Technology*, pages 115–121, New London, CT, March 1995
3. Craig E. Wills. Making a user's information space more visible. In *Proceedings: First Annual Conference on Telecommunications R&D in Massachusetts, Volume I*, pages 75–85, Lowell, MA, October 1994

Other Publications

1. Craig E. Wills. U.S. federal election competitiveness and vote importance. Technical Report WPI-CS-TR-23-04, Computer Science Department, Worcester Polytechnic Institute, July 2023. Available via SSRN.

2. Craig E. Wills. American college influence—a geographical perspective. Technical Report WPI-CS-TR-23-03, Computer Science Department, Worcester Polytechnic Institute, May 2023. Available via SSRN.
3. Craig E. Wills. Self-identified college peer groups and derived rankings. Technical Report WPI-CS-TR-23-01, Computer Science Department, Worcester Polytechnic Institute, January 2023. Available via SSRN.
4. Craig E. Wills and Chayanne Sandoval-Williams. Migration of American college students. Technical Report WPI-CS-TR-22-06, Computer Science Department, Worcester Polytechnic Institute, 2022. Available via SSRN.
5. Craig E. Wills. A data-driven approach to evaluate the worthiness of markets for professional sports franchises. Technical Report WPI-CS-TR-21-02, Computer Science Department, Worcester Polytechnic Institute, February 2021. Available via SSRN.
6. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2022. Technical Report WPI-CS-TR-21-07, Computer Science Department, Worcester Polytechnic Institute, November 2021.
7. Abigail R. Roane, Chaiwat Ekkaewnumchai, Connor W. McNamara, Kyle Richards, Gabor Sarkozy, and Craig E. Wills. A graph-based approach to better sports rankings. Technical Report WPI-CS-TR-19-03, Computer Science Department, Worcester Polytechnic Institute, June 2019.
8. Alan Ritacco and Craig Wills. Peering into the home network. Technical Report WPI-CS-TR-18-02, Computer Science Department, Worcester Polytechnic Institute, April 2018.
9. Craig E. Wills. Geographical connectivity in the United States. Technical Report WPI-CS-TR-17-01, Computer Science Department, Worcester Polytechnic Institute, June 2017. Available via SSRN.
10. Craig E. Wills. Impact of STEM focus on graduation rates in ranking colleges. Technical Report WPI-CS-TR-16-05, Computer Science Department, Worcester Polytechnic Institute, November 2016.
11. Craig E. Wills. A new perspective on United States geography: The closest locations to the most states. Technical Report WPI-CS-TR-16-04, Computer Science Department, Worcester Polytechnic Institute, August 2016. Available via SSRN.
12. Craig E. Wills and Doruk C. Uzunoglu. What ad blockers are (and are not) doing. Technical Report WPI-CS-TR-16-02, Computer Science Department, Worcester Polytechnic Institute, June 2016.
13. Craig E. Wills. The competitiveness of games in professional sports leagues. Technical Report WPI-CS-TR-16-01, Computer Science Department, Worcester Polytechnic Institute, February 2016.

14. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches for 2016. Technical Report WPI-CS-TR-15-03, Computer Science Department, Worcester Polytechnic Institute, November 2015.
15. Craig E. Wills. Analysis of current and future computer science needs via advertised faculty searches. Technical Report WPI-CS-TR-14-06, Computer Science Department, Worcester Polytechnic Institute, November 2014.
16. Craig E. Wills. Analysis of U.S. News graduation rate performance for technological institutions. Technical Report WPI-CS-TR-14-05, Computer Science Department, Worcester Polytechnic Institute, September 2014.
17. Casey Barney, Anthony Caravella, Michael Cullen, Gary Jackson, and Craig E. Wills. Update: Evaluating talent acquisition via the NFL draft. Technical Report WPI-CS-TR-14-03, Computer Science Department, Worcester Polytechnic Institute, May 2014.
18. Casey Barney, Anthony Caravella, Michael Cullen, Gary Jackson, and Craig E. Wills. Evaluating talent acquisition via the NFL draft. Technical Report WPI-CS-TR-13-01, Computer Science Department, Worcester Polytechnic Institute, March 2013.
19. Craig E. Wills and Can Tatar. Understanding what they do with what they know. Technical Report WPI-CS-TR-12-03, Computer Science Department, Worcester Polytechnic Institute, August 2012. This is a longer version of the October 2012 Workshop on Privacy in the Electronic Society paper.
20. Murad Kaplan, Mihajlo Zeljkovic, Mark Claypool, and Craig Wills. JavaScript and Flash overhead in the web browser sandbox. Technical Report WPI-CS-TR-10-14, Computer Science Department, Worcester Polytechnic Institute, April 2012.
21. Balachander Krishnamurthy and Craig E. Wills. Privacy diffusion on the web: A longitudinal perspective (updated graphs), October 2009. Submitted as public comment to Federal Trade Commission Exploring Privacy Roundtable Series.
22. Mark Claypool, Robert Kinicki, and Craig Wills. User-centered network measurement. Technical Report WPI-CS-TR-07-08, Computer Science Department, Worcester Polytechnic Institute, August 2007.
23. Craig E. Wills. Cinderella and the big dance. Technical Report WPI-CS-TR-06-01, Computer Science Department, Worcester Polytechnic Institute, April 2006.
24. Paul J. Timmins and Craig E. Wills. Using future context in personal information retrieval. Technical Report WPI-CS-TR-05-17, Computer Science Department, Worcester Polytechnic Institute, November 2005
25. Hao Shang and Craig E. Wills. Exploiting flow relationships to improve performance of networked applications. Technical Report WPI-CS-TR-04-13, Computer Science Department, Worcester Polytechnic Institute, May 2004.

26. Hao Shang and Craig E. Wills. Using related domain names to improve DNS performance. Technical Report WPI-CS-TR-03-35, Computer Science Department, Worcester Polytechnic Institute, December 2003.
27. Craig E. Wills and Nitin John. A client-based study of clustered and distributed web content. Technical Report WPI-CS-TR-02-27, Computer Science Department, Worcester Polytechnic Institute, October 2002.
28. Balachander Krishnamurthy, Craig Wills, and Yin Zhang. On the use and performance of content distribution networks. Technical Report TD-52AMHL, AT&T Labs – Research, August 2001.
29. Mikhail Mikhailov and Craig E. Wills. Change and relationship-driven content caching, distribution and assembly. Technical Report WPI-CS-TR-01-03, Computer Science Department, Worcester Polytechnic Institute, March 2001.
30. Craig E. Wills, Mikhail Mikhailov, and Hao Shang. N for the price of 1: Bundling web objects for more efficient content delivery. Technical Report WPI-CS-TR-00-26, Computer Science Department, Worcester Polytechnic Institute, November 2000.
31. Craig E. Wills and Hao Shang. The contribution of DNS lookup costs to web object retrieval. Technical Report WPI-CS-TR-00-12, Computer Science Department, Worcester Polytechnic Institute, July 2000.
32. Mikhail Mikhailov and Craig E. Wills. Embedded objects in web pages. Technical Report WPI-CS-TR-00-05, Computer Science Department, Worcester Polytechnic Institute, March 2000.
33. Craig E. Wills and Mikhail Mikhailov. Studying the impact of more complete server information on web caching. Technical Report WPI-CS-TR-99-36, Computer Science Department, Worcester Polytechnic Institute, November 1999.
34. Craig E. Wills and Mikhail Mikhailov. Exploiting object relationships for web caching. Technical Report WPI-CS-TR-99-29, Computer Science Department, Worcester Polytechnic Institute, October 1999.
35. Craig E. Wills and Mikhail Mikhailov. Examining the cacheability of user-requested web resources. Technical Report WPI-CS-TR-99-01, Computer Science Department, Worcester Polytechnic Institute, January 1999.
36. Craig E. Wills and Mikhail Mikhailov. Towards a better understanding of web resources and server responses for improved caching. Technical Report WPI-CS-TR-98-27, Computer Science Department, Worcester Polytechnic Institute, December 1998.
37. Craig E. Wills, Dorothy Deremer, Renee A. McCauley, and Linda Null. Studying the use of peer learning in the introductory computer science curriculum. Technical Report WPI-CS-TR-97-11, Computer Science Department, Worcester Polytechnic Institute, September 1997.

38. Craig E. Wills and Joel Sommers. Prefetching on the web through merger of client and server profiles, June 1997.
39. Joel Sommers and Craig E. Wills. Prefetching on the web using client and server profiles. Technical Report WPI-CS-TR-97-2, Computer Science Department, Worcester Polytechnic Institute, June 1997
40. C.E. Wills, M.T. Murray, and R. Thangarajah. Resource-efficient policies for information transfer in a mobile environment. Technical Report WPI-CS-TR-96-3, Computer Science Department, Worcester Polytechnic Institute, December 1996

Patents

1. B. Krishnamurthy, A. Bender, and C.E. Wills. Protection of personally identifiable information. United States Patent No. 11,003,782 issued May 11, 2021. Continuation of U.S. application Ser. No. 15/631,087 filed Jun. 23, 2017
2. B. Krishnamurthy, A. Bender, and C.E. Wills. Tailored protection of personally identifiable information. United States Patent No. 10,579,804 issued March 3, 2020. Continuation of U.S. application Ser. No. 14/874,493 filed Oct. 5, 2015
3. B. Krishnamurthy, A. Bender, and C.E. Wills. Tailored protection of personally identifiable information. United States Patent No. 9,721,108 issued August 1, 2017. Continuation of U.S. application Ser. No. 12/624,012 filed Nov. 23, 2009
4. B. Krishnamurthy, A. Bender, and C.E. Wills. Tailored protection of personally identifiable information. United States Patent No. 9,172,706 issued October 27, 2015. Provisional patent claim (12/624,012) filed November 23, 2009. Continuation in September 2015
5. B. Krishnamurthy and C.E. Wills. Identifying and remedying secondary privacy leakage. United States Patent No. 8,839,443 issued September 16, 2014. Provisional patent claim (12/288,071) filed October 16, 2008
6. B. Krishnamurthy and C.E. Wills. Method and apparatus for providing web privacy. United States Patent No. 8,601,591 issued December 3, 2013. Provisional patent (12/569,491) filed September 29, 2009
7. B. Krishnamurthy and C.E. Wills. Method for improving web performance by adapting servers based on client cluster characterization. United States Patent No. 7,296,089 issued November 13, 2007. Provisional patent claim (60/346366) filed Nov. 9, 2001. Full claim filed with U.S. Patent and Trademark Office on July 2002
8. B. Krishnamurthy and C.E. Wills. Method and apparatus for cache validation for proxy caches. Patent Number 6,578,113 issued June 10, 2003. Provisional patent claim (60/047,380) filed June 2, 1997. Full claim filed with U.S. Patent and Trademark Office on December 30, 1997

In the News

1. Isha Trivedi. Where do students go to college? a new study looks state by state. *The Chronicle of Higher Education*, June 2, 2022.
2. Alexander C. Kafka. The discipline that is transforming higher ed: The computer-science boom is straining colleges. but it could save some, too. *The Chronicle of Higher Education*, April 15 2020.
3. Jen A. Miller. What does it take to prepare graduates for a new world of work? *EdTech*, May 16 2019.
4. Roberta Kwok. Junior AI researchers are in demand by universities and industry. *Nature*, April 23 2019.
5. William Terdoslavich. Tech industry really needs professors and teaching talent. *Dice*, April 30 2018.
6. Craig Wills. Steps to curb exposure of your data. *Worcester Telegram & Gazette*, April 22 2018.
7. The value of learning to code. *Boston Globe*, August 2 2015. Boston Globe Magazine Education+Careers supplement
8. Ellen O’Leary. Massachusetts schools strive to increase access to coding courses. *Boston.Com*, October 23 2014.
9. Aliya Sternstein. Taking a flier on big data. *Government Executive*, May 28 2013.
10. Shalise Manza Young. Patriots’ draft strategy backed by the numbers. *Boston Globe*, May 12 2013.
11. Better value in 2nd-round picks. *ESPN*, April 26 2013. Story also appeared in Sports Illustrated, Washington Post, Boston Globe, Worcester News Tonight and The Atlantic. Radio Interviews on Chicago, Orlando and Boston Sports Talk Radio shows.
12. Jennifer Toland. NFL draft: WPI study analyzes value of picks. *Worcester Telegram & Gazette*, April 26 2013.
13. Jennifer Valentino-Devries and Jeremy Singer-Vine. They know what you’re shopping for. *Wall Street Journal*, December 7 2012.
14. Office of the Privacy Commissioner of Canada. News release: Popular websites in canada disclosing personal information, September 25 2012.
15. Byron Acohido. Web tracking has become a privacy time bomb. *USA Today*, August 3 2011.
16. Sruthi Krishnan. How your visits to sites are tracked. *The Hindu*, December 25 2010.
17. Ira Flatow. Internet privacy: Who’s tracking you online? *National Public Radio Science Friday*, December 17, 2010.

18. Byron Acohido and Jon Swartz. Do not track could revolutionize online ad industry. *USA Today*, December 13 2010.
19. Geoffrey A. Fowler and Emily Steel. Myspace, apps leak user data. *Wall Street Journal*, October 23 2010.
20. Julia Angwin. The web's new gold mine: Your secrets. *Wall Street Journal*, July 30 2010.
21. Bill Snyder. You are here: Scary new location privacy risks. *CIO*, June 28 2010.
22. Rebecca Myles. Radio interview on wbai evening news. *WBAI, Pacifica Radio 99.5 FM in New York City*, May 26 2010.
23. Michael Hiltzik. Is your privacy secure online? there's no way to tell. *Los Angeles Times*, June 06 2010.
24. Scott Duke Harris. Facebook overhaul simplifies privacy controls. *San Jose Mercury News*, May 27 2010.
25. Rebecca Myles. Radio interview on wbai evening news. *WBAI, Pacifica Radio 99.5 FM in New York City*, May 26 2010.
26. Emily Steel and Jessica E. Vascellaro. Facebook, myspace confront privacy loophole. *Wall Street Journal*, May 21 2010.
27. Lucy Soto. Companies use users' web information to their advantage. *Atlanta Journal-Constitution*, February 12 2010.
28. Erika Morphy. Creepy ways your social media data can be used. *TechNewsWorld*, January 21 2010.
29. Wendy Davis. Social networks may 'leak' personally identifiable information. *MediaPost News*, September 25 2009.
30. Jaikumar Vijayan. Social networking sites leaking personal information to third parties, study warns. *ComputerWorld*, September 23 2009. This syndicated article also appeared in the San Francisco Chronicle, InfoWorld and MacWorld.
31. Peter Eckersley. How online tracking companies know most of what you do online (and what social networks are doing to help them). *Electronic Frontier Foundation Deeplinks Blog*, September 21 2009.
32. Robert Westervelt. Social network privacy study finds identity link to cookies. *SearchSecurity.com*, August 26 2009.
33. Thomas Claburn. Social networks leak personal information. *InformationWeek*, August 24 2009.
34. Miguel Helft. Google is top tracker of surfers in study. *The New York Times Bits Blog*, June 2 2009.

Professional Activities/Honors

1. Recognized as an Academic Advisor with a significant number of academic advisees by Committee on Advising and Students Life. February 2004
2. Recognized as an Academic Advisor with a significant number of academic advisees by Committee on Advising and Students Life. February 2003
3. Selected as an Outstanding Academic Advisor by Committee on Advising and Students Life. March 2000
4. Selected as an honorary member of the WPI Upsilon Pi Epsilon (UPE) computer science honor society. Fall 1995. Quoting from the invitation letter from Scott Salvidio, chapter president, "Your election as an honorary member of Upsilon Pi Epsilon signifies the high regard we have for your work in the field of computer science as a faculty member of WPI. The level of dedication and hard work you exhibit while educating the members of the computer science student body is second to none. Considering the nature of your position in the computer science department we feel that your contributions to the WPI CS community are all the more deserving of special recognition."
5. Member, Association for Computing Machinery (ACM)
6. Member, IEEE Computer Society

Professional Chairs and Editorships

1. Associate Editor, ACM Transactions on Internet Technology, July 2000–July 2009. One of the founding Associate Editors
2. Deputy Program Chair, Ninth International Workshop on Web Caching and Content Delivery, October 2004, Beijing, China
3. Conference Program Committee Co-Chair, 13th International World Wide Web Conference, May 2004, New York City, New York
4. Program Committee Vice-Chair Performance and Reliability, 12th International World Wide Web Conference, May 2003, Budapest, Hungary
5. Panels Chair, Program Committee Member, 11th International World Wide Web Conference, May 2002, Honolulu, Hawaii
6. Faculty Posters Coordinator, ACM SIGCSE 2002, Feb/Mar 2002, Covington, Kentucky, USA
7. Program Committee Deputy Vice-Chair, 10th International World Wide Web Conference, May 2001, Hong Kong

Additional Program Committees

1. Program Committee for ACM/IEEE Hot Topics on Web of Things (HotWot) 2020, November 2020, San Jose USA
2. Intelligent Systems and Infrastructure Track Program Committee of The Web Conference 2020, April 2020, Taipei, Taiwan
3. Program Committee for ACM/IEEE Hot Topics on Web of Things (HotWot) 2019, November 2019, Washington, DC USA
4. Intelligent Systems and Infrastructure Track Program Committee of The Web Conference 2019, May 2019, San Francisco, CA USA
5. Program Committee for ACM/IEEE Hot Topics on Web of Things (HotWot) 2018, October 2018, Bellevue, WA USA
6. Program Committee for HotWeb 2017, October 2017, San Jose, CA USA
7. Program Committee for HotWeb 2016, October 2016, Washington, DC USA
8. Program Committee for HotWeb 2015, November 2015, Washington, DC USA
9. Program Committee for the Web-Based Systems and Applications track of the 33rd IEEE International Conference on Distributed Computing Systems, July 2013, Philadelphia, PA USA
10. Program Committee of the Conference on Web Privacy Measurement, May/June 2012, Berkeley, CA USA
11. Program Committee of the W3C Workshop on Web Tracking and User Privacy, April 2011, Princeton, NJ USA
12. Program Committee of the Performance, Scalability and Availability track for the 2010 International World Wide Web Conference, April 2010, Raleigh, NC USA
13. Program Committee for Workshop on Online Social Networks ACM SIGCOMM Conference, August 2009, Barcelona, Spain
14. Program Committee of the Performance and Scalability track for the 2009 International World Wide Web Conference, April 2009, Madrid, Spain
15. Program Committee of the IEEE International Conference on Self-Adaptive and Self-Organizing Systems, October 2008, Venice, Italy
16. Program Committee of the 2008 Passive and Active Measurement Conference, April 2008, Cleveland, OH USA
17. Program Committee of the Performance and Scalability track for the 2008 International World Wide Web Conference, April 2008, Beijing, China

Significant Academic Service

1. Invited Member, Provost's Program Performance Committee Chaired by Steve Taylor and Debora Jackson, 2020-21
2. Co-Chair, Search Committee Humanities & Arts Department Head, 2018-19
3. Member, WPI Academic Space Committee , 2015-18
4. Faculty Member, WPI Board of Trustees Budget & Finance Committee , 2014-17
5. Member, WPI Strategic Planning Pillar 6 on Enhancing Capacity, 2014-15
6. Member, WPI Committee on Appointments and Promotions, 2010-11. Elected April, 2010
7. Member, WPI Committee on Tenure and Academic Freedom, 2004-08. Chair, 2006-2007. Secretary, 2005-2006. Elected April, 2004
8. Program Committee Member, Computing Research Association Conference, June 2006, Snowbird, Utah
9. Member, WPI Committee on Academic Policy, 2002-04. Elected April, 2002
10. Member, WPI Committee on Academic Policy, 1998-99. Elected April, 1998
11. Chair, WPI Committee on Academic Operations, 1995-96. Elected May, 1995. Elected to committee April, 1993 and served June, 1993–June, 1996

Fellowships and Grants

1. Craig Shue, Craig E. Wills, Robert Walls, and Lorenzo De Carli. Cybercorps SFS renewal: Supporting the federal government workforce, January 1, 2021 – December 31, 2025. National Science Foundation Scholarships for Service, 1941415. \$4,836,782. Awarded August 2020
2. Suzanne Mello-Stark and Craig E. Wills. GenCyber 2018 WPI Cybersecurity Camp, June 1, 2018 – August 31, 2018. National Science Foundation Scholarships for Service, 1503742. \$99,428. Awarded April 2018
3. Suzanne Mello-Stark and Craig E. Wills. WPI GenCyber student summer camp, June 1, 2017 – August 31, 2017. National Science Foundation Scholarships for Service, 1503742. \$99,687. Submitted November 2016
4. Suzanne Mello-Stark and Craig E. Wills. 2016 GenCyber, June 1, 2016 – August 31, 2016. National Science Foundation Scholarships for Service, 1503742. \$99,769. Submitted November 2015. Awarded April 2016
5. Kathryn Fisler, Craig Shue, Susan Landau, and Craig E. Wills. Scholarship track: Scholarships for service at WPI, January 1, 2015 – December 31, 2019. National Science Foundation Scholarships for Service, 1503742. \$2,831,124. Submitted October 2014. Awarded December 2014 for \$4.4M

6. Mark Claypool and Craig E. Wills. Measuring DNS performance, January, 2011–October, 2011. Dynamic Network Services, Inc. \$55,297
7. Craig E. Wills, Mark Claypool, James Doyle, and Matthew Ward. MRI-R2: Development of a user-centered network measurement platform, May 1, 2010 – April 30, 2013. National Science Foundation Major Research Instrumentation Program Recovery and Reinvestment. 0959441. \$391,582
8. Craig E. Wills. CSR-PDOS: virtual machines meet application clusters: A highly responsive global utility computing platform for internet applications, May, 2009–August, 2009. Research Experience for Undergraduates (REU) Supplement to collaborative proposal with Case-Western Reserve University. 0937144. \$8,000
9. Craig E. Wills, Mark Claypool, and Robert Kinicki. A dual-core experimental systems laboratory, August, 2007. Equipment Donation from Intel Corporation. \$23,959
10. Michael Rabinovich and Craig E. Wills. CSR-PDOS: virtual machines meet application clusters: A highly responsive global utility computing platform for internet applications, August, 2006–July, 2009. National Science Foundation CSR PDOS. 0615079. Collaborative proposal with Case-Western Reserve University. WPI portion \$238,380
11. Mark L. Claypool, Robert E. Kinicki, and Craig E. Wills. Research resources for network application studies, July, 2004–August, 2007. National Science Foundation CNS CISE Research Resources. 0423362. \$39,203
12. Craig E. Wills. Exploiting object relationships for more deterministic management of distributed objects, Sept, 2000–Aug, 2003. National Science Foundation Operating Systems and Compilers Program of the CCR Division in the CISE Directorate. 9988250. \$66,141
13. Mark L. Claypool, David Finkel, and Craig E. Wills. Teaching systems courses with an open source laboratory, June, 2000–May, 2003. National Science Foundation Course, Curriculum and Laboratory Improvement Grant DUE9980803. \$69,912
14. Craig E. Wills. Impact of electronic commerce on the internet infrastructure, September, 1999–May, 2000. Arrowpoint Communications, Inc. \$33,453
15. C.E. Wills, D. Finkel, G.T. Heineman, R.E. Kinicki, and M.O. Ward. The webware, interfaces and networking experimental laboratory, June, 1997–May, 1999 (extended to May, 2000). National Science Foundation Instrumentation and Laboratory Improvement Grant DUE9751132. \$44,256
16. C.E. Wills. Application of peer learning to the introductory computer science curriculum, June, 1996–May, 1998 (extended to May, 2000). National Science Foundation Undergraduate Faculty Enhancement Grant DUE9554706. \$56,521
17. D.C. Brown, C.E. Wills, D. Finkel, N.I. Hachem, R.E. Kinicki, and M.O. Ward. The enhancement of digital's technology exchange program, June, 1995–May, 1996. Digital Equipment Corporation. \$88,960

18. D.C. Brown and C.E. Wills. Continuation of tennis project: Computer network ease of service evaluation, January, 1995–December, 1995. Digital Equipment Corporation. \$83,770
19. C.E. Wills. A networked, platform-independent audio support system, October, 1994–May, 1995. Vicorp Interactive Systems. \$19,004
20. D.C. Brown and C.E. Wills. Computer network serviceability evaluation system, September, 1993–August, 1994. Digital Equipment Corporation. \$68,784