IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent of:	Vipul Pabari	Attorney Docket No. 42342-0157IP1
U.S. Patent No.:	12,079,667	
Issue Date:	September 3, 2024	
Appl. Serial No.:	18/135,007	
Filing Date:	April 14, 2023	
Title:	VIRTUAL SYSTEMS MANAGEMENT	

DECLARATION OF EXPERT

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

Date: May 7, 2026

By: _____
Michael Mitzenmacher

MICROSOFT 1003

Table of Contents

I. ASSIGNMENT....8
II. QUALIFICATIONS AND BACKGROUND INFORMATION....8
III. OVERVIEW OF CONCLUSIONS FORMED....14
IV. LEVEL OF ORDINARY SKILL IN THE ART....14
V. LEGAL STANDARDS....15
A. Terminology....15
B. Legal Standards....16
C. Anticipation....17
D. Obviousness....17
VI. MATERIALS CONSIDERED....21
VII. TERMINOLOGY....22
A. "Virtual Asset"....22
B. "Resources"/"Physical Resources"....23
C. "Software Instantiations of Computer Systems"....23
D. "Allocation of Data"....24
E. "Control Computer"....24
F. "Logical"....25
G. "Middleware"....26
H. "Configuration Management"....26
I. "Application Routing for Virtual Assets"....27
VIII. THE '667 PATENT....29
A. Overview of the '667 Patent....29
B. Prosecution History of the '667 Patent....30
IX. OVERVIEW OF THE PRIOR ART....31
A. Hashimoto (EX1005)....31
B. Lubbers (EX1006)....36
C. Hirschfeld (EX1007)....40
D. Khandekar (EX1012)....43
X. OVERVIEW OF THE PRIOR ART COMBINATIONS....45
A. The Combination of Hashimoto and Lubbers....45
B. The Combination of Hashimoto, Lubbers, and Hirschfeld....50
C. The Combination of Hashimoto, Lubbers, and Khandekar....54
XI. THE PRIOR ART MAKES CLAIMS 1-20 OBVIOUS....55
A. Ground 1: Claims 1-5, 7, and 8 Are Obvious Over The Combination of Hashimoto and Lubbers....55

1. Claim 1 56
2. Claim 2 74
3. Claim 3 76
4. Claim 4 77
5. Claim 5 78
6. Claim 7 80
7. Claim 8 82

B. Ground 2: Claims 9-20 Are Obvioius Over The Combination of Hashimoto, Lubbers, and Hirschfeld 85
1. Claim 9 85
2. Claim 10 94
3. Claim 11 94
4. Claim 12 99
5. Claim 13 99
6. Claim 14 100
7. Claim 15 100
8. Claim 16 101
9. Claim 17 101
10. Claim 18 102
11. Claim 19 106
12. Claim 20 106

C. Ground 3: Claims 6 Are Obvioius Over The Combination of Hashimoto, Lubbers, and Khandekar 106
1. Claim 6 106

XII. CONCLUSION 108

LISTING OF CHALLENGED CLAIMS

Claim 1	
[1pre]	A method for data storage across one or more virtualization system resources over a network, the method comprising:
[1.1.a]	implementing a virtualization system that includes a control center application, one or more physical resources, and one or more virtual assets on each of the physical resources,
[1.1.b]	wherein the virtual assets are managed by virtualization layer instances running on each of the physical resources, and
[1.1.c]	wherein at least two of the virtual assets on one or more of the physical resources include virtual storage assets;
[1.2]	obtaining a set of data to be stored the virtual storage assets on any of the physical resources;
[1.3]	selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual storage assets; and
[1.4]	causing storage of the set of data at the selected virtual storage assets on any of the physical resources.
Claim 2	
[2]	The method of claim 1, wherein the at least one virtual storage asset is virtualized via storage virtualization layers running on each of the physical resources.
Claim 3	
[3]	The method of claim 1, wherein at least two virtual assets each comprise software instantiations of computer systems.
Claim 4	
[4]	The method of claim 2, wherein at least one of the virtual assets comprise middleware applications running on each of the physical resources to access the virtual storage assets via the storage virtualization layer.
Claim 5	
[5]	The method of claim 1, wherein the control center application is configured to perform configuration management or forecasting services via the virtualization layer instances.
Claim 6	

[6]	The method of claim 5, wherein the configuration management includes application routing for virtual assets or implementation of application rules.
Claim 7	
[7]	The method of claim 1, wherein the at least one virtual storage asset is a logical unit with a logical unit number (LUN).
Claim 8	
[8.1]	The method of claim 1, further comprising: obtaining a request from a user device to obtain at least a portion of the set of stored data;
[8.2]	retrieving at least the portion of the set of stored data from the virtual storage asset; and
[8.3]	transmitting at least the portion of the set of stored data to the user device via a network.
Claim 9	
[9pre]	A virtualization system comprising:
[9.1]	a control computer running a control center application; and
[9.2.a]	one or more physical resources, wherein one or more physical resources are configured to execute one or more virtual assets,
[9.2.b]	the virtual assets managed by virtualization layer instances on each of the physical resources,
[9.2.c]	wherein at least two of the virtual assets on the one or more of the physical resources include virtual storage assets,
[9.3]	wherein the control center application is configured to: obtain a request from a user device for data stored at any of the virtual storage assets;
[9.4]	retrieve the data from one or more virtual storage assets that are storing the data; and
[9.5]	transmit the retrieved data to the user device.
Claim 10	
[10]	The virtualization system of claim 9, wherein the virtual storage assets are virtualized via a storage virtualization layer.

Claim 11	
[11.1]	The virtualization system of claim 9, wherein the virtualization system is further configured to: obtain the data to be stored at any of the virtual storage assets; and
[11.2]	store the data across any of the virtual storage assets running on the physical resources.
Claim 12	
[12]	The virtualization system of claim 9, wherein each of the one or more physical resources execute at least two virtual assets that each comprise software instantiations of computer systems.
Claim 13	
[13]	The virtualization system of claim 9, wherein at least one of the virtual assets comprise middleware applications accessing the virtual storage assets via the storage virtualization layer.
Claim 14	
[14]	The virtualization system of claim 9, wherein the control center application is configured to perform configuration management or forecasting via the virtualization layer instances.
Claim 15	
[15]	The virtualization system of claim 9, wherein the virtual storage assets comprise databases.
Claim 16	
[16]	The virtualization system of claim 9, the virtual storage assets are logical units with a logical unit number (LUN).
Claim 17	
[17]	The virtualization system of claim 9, wherein the virtualization system is further configured to receive, over a network by a user, one or more instructions to configure the virtual storage assets.
Claim 18	
[18pre]	A method for data storage on virtualized resources, the method comprising:
[18.1.a]	receiving a request to store data, at a control center application in communication with one or more physical resources that are part of a

	virtualization system that includes the control center application, the one or more physical resources, and one or more virtual assets on each of the physical resources, and
[18.1.b]	wherein at least two of the virtual assets on one or more of the physical resources include virtual data storage assets;
[18.2]	selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual data storage assets; and
[18.3]	causing storage of the data, by the control center application, at the selected virtual data storage assets on each of the one or more physical resources,
[18.4]	wherein the virtual data storage assets are managed by storage virtualization layer instances on each of the physical resources.
Claim 19	
[19]	The method of claim 18, wherein one or more virtual assets are configured to run on each of the physical resources, wherein the virtual assets are managed by virtualization layer instances running on each of the physical resources.
Claim 20	
[20]	The method of claim 19, wherein at least two of the virtual assets comprise software instantiations of computer systems, and wherein the virtual storage assets comprise databases.

I. ASSIGNMENT

1. I have been retained on behalf of Microsoft Corporation ("Microsoft") to offer technical opinions related to 12,079,667 ("the '667 patent") (EX1001). I understand that Microsoft is requesting that the Patent Trial and Appeal Board ("PTAB" or "Board") to institute an *inter partes* review ("IPR") proceeding of the '667 patent.

2. I have been asked to provide my independent analysis of the '667 patent based on the prior art publications cited in this declaration.

3. I received no compensation for this declaration beyond my normal hourly compensation based on my time actually spent analyzing the '667 patent, the prior art publications cited below, and issues related thereto, and I will not receive any added compensation based on the outcome of any IPR or other proceeding involving the '667 patent.

II. QUALIFICATIONS AND BACKGROUND INFORMATION

4. I am over the age of 18 and am competent to write this declaration. I have personal knowledge, or have developed knowledge of these technologies based upon education, training, or experience, of the matters set forth herein.

5. My qualifications are summarized here and explained in more detail in my curriculum vitae, which is attached as Appendix A to this report. Appendix A also includes a list of my publications and the cases in which I have testified at

deposition, hearing, or trial within the past four years.

6. I am currently employed as a Professor of Computer Science at Harvard University. Specifically, I am the Thomas J. Watson, Sr. Professor of Computer Science in the School of Engineering and Applied Sciences. I joined the faculty of Harvard as an Assistant Professor in January 1999. I was promoted to Associate Professor in 2002 and to Professor in 2005. In 2010, I began a three-year term as Area Dean, which is essentially equivalent to what other schools call Department Chair of Computer Science, and held that position through June 2013. I also served as the Area Co-Chair for Computer Science from July 1, 2018, to June 30, 2019. My work address is 150 Western Ave, Allston, MA 02134.

7. I received my undergraduate degree in Mathematics and Computer Science from Harvard College in 1991. I received a Certificate of Advanced Study in Mathematics from Cambridge University in 1992. I received a Ph.D. in Computer Science from the University of California at Berkeley in 1996. From August 1996 to January 1999, I was employed as a Research Scientist at Digital Systems Research Center.

8. I have published over 250 research papers in computer science conferences and journals. My general area of research includes algorithms and data structures, as well as the analysis of probabilistic processes. I am listed as an inventor or co-inventor on 19 issued patents and am the co-author of a textbook,

entitled *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, published by Cambridge University Press. I am a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) and a Fellow of the Association for Computing Machinery (ACM). In addition, I have consulted for several companies while I have been a professor, such as Microsoft, Mitsubishi, Akamai, ITA software, Fluent Mobile, Adverplex, and e-Harmony.

9. I teach courses on various computer science topics to undergraduate and graduate students at Harvard. I have regularly taught the undergraduate class on algorithms and data structures. I have also regularly taught graduate level classes involving algorithms and data structures, that include topics related to machine learning. For example, in both courses I have covered the area of learning-augmented algorithms, where algorithms use predictions to improve their performance.

10. I have both done research in the area of electronic commerce, as well as worked for companies related to electronic commerce. For example, I have written papers regarding how the use of Groupon can affect ratings (such as Yelp ratings) for businesses, on the issue of information asymmetry in pay-per-bid auction sites, and on algorithms for finding statistically significant frequent item sets. Also, as examples, I have served on the scientific advisory board for e-harmony (a dating service with millions of users), and worked at companies such

as Google, ITA software (which was bought by Google and serves as the basis for Google flights) and Adverplex/Cogolabs (which works in the space of web and Internet advertising).

11. Much of my work involves issues relating to networking and communications. I regularly serve on program committees for conferences in networking, algorithms, and communication. For example, I have served on the program committee multiple times for the SIGCOMM conference, which is described on the conference homepage as follows: "SIGCOMM is the flagship annual conference of the ACM Special Interest Group on Data Communication (SIGCOMM) on the applications, technologies, architectures, and protocols for computer communication." https://sigcomm.org/events/sigcomm-conference. Similarly, I have served several times on the Program Committee for NSDI, the USENIX Symposium on Networked Systems Design and Implementation. I have written papers on networking that have been published in the IEEE/ACM Transactions on Computer Networking, the SIGCOMM conference, the INFOCOM conference, and other major venues for networking research. My graduate course entitled "Algorithms at the end of the wire" covers many subjects at the intersection of networking and algorithms.

12. My curriculum vitae, included as an appendix to this declaration, includes a list of publications on which I am a named author. It contains further

details regarding my experience, education, publications, and other qualifications to render an expert opinion in connection with this proceeding.

13. In writing this Declaration, I have considered the following: my own knowledge and experience, including my work experience in the fields of networking and communications; my experience in teaching those subjects; and my experience in working with others involved in those fields. In addition, I have analyzed the following publications and materials, in addition to other materials I cite in my declaration:

14. I have reviewed the ’667 patent (EX1001), and relevant excerpts of the prosecution history (EX1002) of the ’667 patent. Among various textbooks, documents, and publications, I have also reviewed the following prior art references:

- U.S. Patent Publication No. 2006/0174087 to Akiyoshi Hashimoto *et al.* (“Hashimoto”) (EX1005)
- U.S. Patent Publication No. 2003/0084241 to Clark E. Lubbers *et al.* (“Lubbers”) (EX1006)
- U.S. Patent Publication No. 2003/0051021 to Robert A. Hirschfeld *et al.* (“Hirschfeld”) (EX1007)
- U.S. Pub. No. 2006/0294519 to Naoya Hattori *et al.* (“Hattori”) (EX1011)
- U.S. Patent No. 7,577,722 to Dilip Khandekar *et al.* (“Khandekar”)

(EX1012)

- U.S. Pat. No. 8,335,794 to John Goodson *et al.* ("Goodson") (EX1013)
- U.S. Pub. No. 2006/0136667 to Steven Shultz *et al.* ("Shultz") (EX1014)
- U.S. Pub. No. 2003/0225801 to Murthy Devarakonda *et al.* ("Devarakonda") (EX1015)

15. Counsel (Fish & Richardson) has informed me that I should consider these materials through the lens of one of ordinary skill in the art related to the '667 patent at the time of the earliest possible priority date of the '667 patent, and I have done so during my review of these materials. I have been informed by Counsel to use the date of August 15, 2005 as the Priority Date.

16. I have no financial interest in the outcome of this proceeding. I am being compensated for my work as an expert on an hourly basis. My compensation is not dependent on the outcome of these proceedings or the content of my opinions.

17. My opinions, as explained below, are based on my education, experience, and expertise in the fields relating to the '667 Patent. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the art as of the earliest possible priority date. Any figures that appear within this document have been prepared with the assistance of Counsel and reflect my understanding of the '667 Patent and the prior art discussed below.

III. OVERVIEW OF CONCLUSIONS FORMED

18. This declaration explains the conclusions that I have formed based on my analysis of all of the considered materials and my personal experience. To summarize those conclusions, based upon my knowledge and experience and my review of the prior art references listed above, I believe that:

- Ground 1: Hashimoto-Lubbers renders claims 1-5, 7, and 8 obvious
- Ground 2: Hashimoto-Lubbers-Hirschfeld renders claims 9-20 obvious
- Ground 3: Hashimoto-Lubbers-Khandekar renders claim 6 obvious

19. In support of these conclusions, I provide an overview of the references in Section IX and more detailed comments regarding the obviousness of claims 1-20 ("the Challenged Claims") of the '667 patent in Sections X and XI.

IV. LEVEL OF ORDINARY SKILL IN THE ART

20. I have been informed that a person of ordinary skill in the art ("POSITA") is a hypothetical person who is presumed to have the skill and experience of an ordinary worker in the field at the time of the alleged invention. Because I do not know at what date the invention as claimed was made, if ever, I have used the Critical Date of the '667 patent as the point in time for claim

interpretation purposes. My opinion does not change if the invention date is earlier.

21. Based on my knowledge and experience in the field and my review of the '667 patent and file history, I believe that a person of ordinary skill in the art in this matter would have had a combination of (1) at least a bachelor's degree in computer science, computer engineering, electrical engineering, or a related field, and (2) at least two years of industry experience with virtualization systems. Additional education could substitute for professional experience, and professional experience could substitute for education.

22. Based on my experiences, I have a good understanding of the capabilities of a POSITA. Indeed, I have taught, participated in organizations, and worked closely with many such persons over the course of my career. Based on my knowledge, skill, and experience, I have an understanding of the capabilities of a POSITA. For example, from my industry experience, I am familiar with what a POSITA would have known and found predictable in the art. From teaching and supervising my post-graduate students, I also have an understanding of the knowledge that a person with this academic experience possesses. Furthermore, I possess those capabilities myself.

V. LEGAL STANDARDS

A. Terminology

23. I have been informed by Counsel and understand that the best

indicator of claim meaning is its usage in the context of the patent specification as understood by one of ordinary skill. I further understand that the words of the claims should be given their plain meaning unless that meaning is inconsistent with the patent specification or the patent's history of examination before the Patent Office. Counsel has also informed me, and I understand that, the words of the claims should be interpreted as they would have been interpreted by one of ordinary skill at the time of the invention was made (not today). The '667 patent was filed April 14, 2023, and claims priority to August 15, 2005, which I refer to as the "Critical Date" in this proceeding. Because I do not know at what date the invention as claimed was made, if ever, I have used the Critical Date of the '667 patent as the point in time for claim interpretation purposes. My opinion does not change if the invention date is earlier.

B. Legal Standards

24. I have been informed by Counsel and understand that documents and materials that qualify as prior art can render a patent claim unpatentable as being anticipated or obvious.

25. I am informed by Counsel and understand that all prior art references are to be looked at from the viewpoint of a person of ordinary skill in the art at the time of the invention, and that this viewpoint prevents one from using his or her own insight or hindsight in deciding whether a claim is anticipated or rendered

obvious.

C. Anticipation

26. I understand that patents or printed publications that qualify as prior art can be used to invalidate a patent claim as anticipated or as obvious.

27. 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.

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

D. Obviousness

29. 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 POSITA.

30. I have been informed by Counsel and understand that a claim is unpatentable for obviousness and that obviousness may be based upon a combination of references. I am informed by Counsel and understand that the combination of familiar elements according to known methods is likely to be

obvious when it does no more than yield predictable results. However, I am informed by Counsel and understand that a patent claim composed of several elements is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art.

31. I am informed by Counsel and understand that when a patented invention is a combination of known elements, a court determines whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue by considering the teachings of prior art references, the effects of demands known to people working in the field or present in the marketplace, and the background knowledge possessed by a person having ordinary skill in the art.

32. I am informed by Counsel and understand that a patent claim composed of several limitations is not proved obvious merely by demonstrating that each of its limitations was independently known in the prior art. I am informed by Counsel and understand that identifying a reason those elements would be combined can be important because inventions in many instances rely upon building blocks long since uncovered, and claimed discoveries almost of necessity will be combinations of what, in some sense, is already known. I am informed by Counsel and understand that it is improper to use hindsight in an obviousness analysis, and that a patent's claims should not be used as a "roadmap."

33. I am informed by Counsel and understand that an obviousness inquiry requires consideration of the following factors: (1) the scope and content of the prior art, (2) the differences between the prior art and the claims, (3) the level of ordinary skill in the art, and (4) any so called "secondary considerations" of non-obviousness, which include: (i) "long felt need" for the claimed invention, (ii) commercial success attributable to the claimed invention, (iii) unexpected results of the claimed invention, and (iv) "copying" of the claimed invention by others.

34. I have been informed by Counsel and understand that an obviousness evaluation can be based on a single reference or 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 that the nexus linking two or more prior art references is sometimes simple common sense. I have been informed by Counsel and 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.

35. I have been informed by Counsel and understand that if a technique has been used to improve one device, and a person of ordinary skill at the time of invention would have recognized 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.

36. I have been informed by Counsel and 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 have been informed by Counsel and understand that a person of ordinary skill looking to overcome a problem will often be able to fit together the teachings of multiple prior art references. I have been informed by Counsel and understand that obviousness analysis therefore takes into account the inferences and creative steps that a person of ordinary skill would have employed at the time of invention.

37. I have been informed by Counsel and understand that a proper obviousness analysis focuses on what was known or obvious to a person of ordinary skill at the time of invention, 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.

38. I have been informed by Counsel and 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.

39. In sum, my understanding is that prior art teachings are properly

combined where one of ordinary skill 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.

40. I have been informed by Counsel and understand that in an *inter partes* review, "the petitioner shall have the burden of proving a proposition of unpatentability," including a proposition of obviousness, "by a preponderance of the evidence." 35 U.S.C. §316(e).

VI. MATERIALS CONSIDERED

41. My analysis and conclusions set forth in this declaration are based on my educational background and experiences in the field (*see* Section II). Based on my knowledge and experience, I believe that I am considered to be an expert in the field. Also, based on my knowledge and experience, I understand and know of the capabilities of persons of ordinary skill in the field around the Critical Date.

42. As part of my independent analysis for this declaration, I have considered the following: the background knowledge/technologies that were commonly known to persons of ordinary skill in this art during the time before the Critical Date; my own knowledge and experiences gained from my work

experience in the field of the '667 patent and related disciplines; and my experience in working with others involved in this field and related disciplines.

43. In addition, I have analyzed the publications and materials listed above.

44. My analysis and conclusions set forth in this declaration are based on the perspective of a POSITA.

VII. TERMINOLOGY

45. I believe the prior art grounds I describe below would render the Challenged Claims obvious under any reasonable claim constructions, including plain and ordinary meaning. For clarity, though, counsel provided the following constructions, which I believe are supported by the disclosure of the '667 Patent, and therefore demonstrate subject matter that falls within the metes and bounds of these terms. I reserve the right to respond to any constructions offered by Patent Owner or adopted by the Board.

A. "Virtual Asset"

46. The '667 Patent specifies that "virtual devices that exist only as software instantiations of equipment objects in a virtual environment will be referred to as virtual assets." EX1001, 6:43-47. The '667 Patent further explains that "[t]he resource-asset dichotomy will be maintained throughout this discussion." EX1001, 6:47-48. Accordingly, accounting for context offered by the

'667 Patent, a POSITA would have understood that the claimed "virtual asset" encompasses "software instantiations of equipment objects that exist only within a virtual environment." EX1001, 6:43-48.

B. "Resources"/"Physical Resources"

47. The '667 Patent specifies that, "[i]n this description, physical devices such as computers, printers, routers, and other "boxes" will be referred to as resources." EX1001, 6:43-47. The '667 Patent further explains that "[t]he resource-asset dichotomy will be maintained throughout this discussion." EX1001, 6:47-48. Accordingly, accounting for context offered by the '667 Patent, a POSITA would have understood that the claimed "resources" and "physical resources" encompasses "physical devices on the network, such as computers, printers, routers, and other tangible hardware." EX1001, 6:43-48.

C. "Software Instantiations of Computer Systems"

48. The '667 Patent states that "[virtual] assets…provide the functional equivalent of desktops (user interfaces), operating systems, applications, servers, data bases, and the like." EX1001, 3:17-21. In FIG. 3, the virtual assets are illustrated as two virtual computers, or machines, a first Virtual Machine 328 and a second Virtual Machine 330." EX1001, 10:10-12. "Each virtual machine can include an operating system (OS) for the virtual machine, computer applications programs executing on the virtual machine OS, and virtual hardware assets for the

virtual machine." EX1001, 10:13-16. Accordingly, accounting for context offered by the '667 Patent, a POSITA would have understood that the claimed "software instantiations of computer systems" encompasses "software that can virtually perform functions of physical devices that can connect to computers."

D. "Allocation of Data"

49. The '667 Patent states that "FIG. 12 is a block diagram representation of how storage in physical servers can be allocated by the virtual server using the system management of the present invention." EX1001, 20:58-60. FIG. 12 "compris[es] an association or mapping between actual Physical Resources comprising a collection of disk drives or storage devices 1204 to a collection of Logical Units 1206, such as network virtual storage drives, in accordance with a set of application critical Rules 1208." EX1001, 20:60-67. The '667 Patent further describes that "[t]he configuration of the Logical Units 1206 and their corresponding virtual drives will be governed by the set of Rules 1208." EX1001, 21:7-9. Accordingly, accounting for context offered by the '667 Patent, a POSITA would have understood that the claimed "allocation of data" encompasses "the mapping of physical storage capacity to logical storage units according to specific application critical rules."

E. "Control Computer"

50. The '667 Patent states that "[t]he Control Center computers are

represented in FIG. 1 as physical devices rather than virtual assets..” EX1001, 6:35-37. “Each Control Center computer 110, 112 that is equipped with the virtualization management application described herein will have management of one or more physical assets within a domain or other subnet arrangement associated with the respective Control Center.” EX1001, 7:13-16. Accordingly, accounting for context offered by the ’667 Patent, a POSITA would have understood that the term “control computer” at least encompasses a “computing device that is separate and distinct from the physical host machines being managed and is not itself deployed as a virtual asset.”

F. “Logical”

51. The ’667 Patent states that “FIG. 12 depicts the configuration management 120 of the network computer system 100 (FIG. 1) as comprising an association or mapping between actual Physical Resources comprising a collection of disk drives or storage devices 1204 to a collection of Logical Units 1206, such as network virtual storage drives, in accordance with a set of application critical Rules 1208.” EX1001, 20:60-67. “That is, the storage servers (physical resources comprising disks) are mapped to virtual assets (logical units).” EX1001, 20:67-21:2. Accordingly, accounting for context offered by the ’667 Patent, a POSITA would have understood that the claimed “logical” encompasses “software-defined abstraction that requires mapping to underlying physical hardware.”

G. "Middleware"

52. The '667 Patent provides multiple examples of "middleware." In one example, the '667 Patent states that "[a] Virtual Platform Abstraction Layer 416 permits communications between the host machine 302 resources and a set of multiple middleware adapters 418." EX1001, 11:48-52. "The middleware adapters communicate with the virtualization environments of the host machine 302." EX1001, 14:39-41.

53. In another example, the '667 Patent describes that "applications can be supported by middleware software through an application virtualization layer 616." EX1001, 16:63-65. "The middleware applications 618, 620, 622 can access data stores through a storage virtualization layer 624." EX1001, 16:67-17:1.

54. Accordingly, accounting for context offered by the '667 Patent, a POSITA would have understood that the claimed "middleware" encompasses "software that acts as a bridge between an operating system or database and applications, especially on a network." *See also* EX1010.

H. "Configuration Management"

55. The '667 Patent states that "the Control Center application can perform configuration management and forecasting services for the host computer and the available virtual assets through each of the various virtualization layers." EX1001, 17:3-8. "FIG. 11 is a block diagram representation of a server

virtualization configuration management scheme that can be implemented from a computer such as the Control Center computers 110, 112 illustrated in FIG. 1." EX1001, 20:32-35. "FIG. 11 depicts the virtualization management 1102 of the network computer system 100 (FIG. 1) as comprising an association or mapping between the actual Physical Resources, such as a collection of computer hardware machines 1104, to a collection of Virtual Machine Assets 1106, or applications, in accordance with a set of application critical Rules 1108." EX1001, 20:35-41. "If configuration changes are called for,…then the Control Center issues commands to the virtualization layer software to implement the desired configuration changes." EX1001, 23:16-19. The '667 Patent describes various management functions, including "Create VM," "Start VM," "Create Virtual Disk," "Consolidate_VM," "Configure CPU," "Configure disk," Configure memory," "configure network," etc. *See* EX1001, 15:42-16:44. Accordingly, accounting for context offered by the '667 Patent, a POSITA would have understood that the claimed "configuration management" encompasses "the creation, modification, and management of associations or mappings between physical resources and virtual assets, performed by the control center application in accordance with a defined set of application rules.

I. "Application Routing for Virtual Assets"

56. In the '667 Patent, "[t]he Control Center can [] provide Adaptive

Application Routing (AAR) for virtual assets based on [Application Criticality Rules] ACRs." EX1001, 18:3-4. "The ACRs allow a user to specify business rules for controlling priority settings for applications." EX1001, 18:8-9. To this end, the Control Center includes a "Dynamic Application Router." EX1001, 7:51-55, 8:54-9:7. "This component can be used to move applications to different virtual assets as global business need changes, provide dynamic allocation of global virtual assets for workload optimization, enable zero down time in upgrading virtual environments, provide an optimizer for balancing resource and asset inventory against needs, a scheduler to auto schedule actions and provide reports, and provide alarms and triggers to notify users of out of balance inventory items." EX1001, 8:59-67. Further, this component "can compare business policy against real time asset usage and then make load balance changes." EX1001, 8:67-9:7. In one example, "if a rule specifies that a particular application has a high priority at a particular time of day, then a respective one of the Control Center computers 110, 112 (FIG. 1) will adjust the deployment of applications among the Virtual Machines 1106 under its management at the appropriate times." Accordingly, accounting for context offered by the '667 Patent, a POSITA would have understood that the claimed "application routing for virtual assets" encompasses "the dynamic deployment, movement, or allocation of applications to or between virtual assets based on business policies, workload conditions, or

application criticality rules, as a function of configuration management performed by the control center application."

VIII. THE '667 PATENT

A. Overview of the '667 Patent

57. The '667 Patent describes "methods and apparatus for management of one or more virtual environments regardless of any underlying central processing unit (CPU) specification and regardless of any underlying operating system (OS) or virtualization environment." EX1001, 3:53-57. However, the claims of the '667 Patent make no mention of CPU specification or underlying OS. *See* EX1001, 24:10-26:30. This contrasts with claims in the earliest patent granted in the family of which the '667 Patent is a member, U.S. Patent No. 8,799,431. *See* EX1008, 23:5-26:38; EX1001, p. 2 (listing '431 Patent). For example, independent claim 20 of the '431 Patent includes the limitation that "functions are carried out without regard to processors, operating systems, virtualization platforms, and application software of the virtualization environment." EX1008, 25:5-7.

58. Rather, the claims of the '667 Patent relate to simple storage and retrieval of data from virtualized storage assets in a straightforward virtualization system. *See* EX1001, 24:10-26:30. For example, claim 1 of the '667 Patent is reproduced below:

> **1**. A method for data storage across one or more virtualization system resources over a network, the method comprising:
>
> implementing a virtualization system that includes a control center application, one or more physical resources, and one or more virtual assets on each of the physical resources, wherein the virtual assets are managed by virtualization layer instances running on each of the physical resources, and wherein at least two of the virtual assets on one or more of the physical resources include virtual storage assets;
>
> obtaining a set of data to be stored the virtual storage assets on any of the physical resources;
>
> selecting one or more of the virtual assets for storage of the set of data based on an allocation of data on the virtual storage assets; and
>
> causing storage of the set of data at the selected virtual storage assets on any of the physical resources.

EX1001, 24:10-27. Unsurprisingly, these basic storage features were well known in the art.

B. Prosecution History of the '667 Patent

59. The '667 Patent was the most recent patent issued in a family that includes four earlier granted patents. *See* EX1001, pp. 1-2 (listing four prior patents under "Related U.S. Application Data"). The '667 Patent in particular received limited examination, with only a single Office Action issued before allowance. *See* EX1002, 83-95. In this Office Action, the Examiner provided a double-patenting rejection based on prior patents in the family, and an obviousness rejection based on U.S. Publication Nos. 2005/0044301 ("Vasilevski") and 2006/0149842 ("Dawson"). *See id.* In response, the Applicant amended independent claims 1 and 18 to recite the "selecting" limitation (Limitations [1.3]

and [18.2]) and argued this was not shown in the prior art. *See* EX1002, 75-82. Notably, Applicant did not amend independent claim 9 to recite a similar "selecting" limitation, and yet argued that the prior art "does not teach or suggest the limitations of independent claim 9…for at least reasons similarly to that provided with respect to claim 1." *See id.*

60. In allowing the '667 Patent, the Examiner indicated that limitations [1.1] and [1.3] were not in the closest prior art he identified. *See* EX1002, 17. However, independent claim 9 does not even include limitation [1.3] or any similar limitations. *See* EX1001, 24:55-25:3.

IX. OVERVIEW OF THE PRIOR ART

A. Hashimoto (EX1005)

61. Hashimoto "relates to a computer system, and more particularly to a technique of logically partitioning a computer system including a storage system." EX1005, [0002]. Hashimoto teaches that "[w]ith conventional techniques of logically partitioning computers, resources in a computer, such as processors and memory, are logically partitioned and allocated to individual virtual machines." EX1005, [0010]. "Similarly, conventional techniques of logically partitioning storage systems logically partition resources in a storage system, such as cache memory and disk drives, and allocate the partitioned resources to individual virtual storage systems." EX1005, [0010]. "With this technology, different operating

systems can be freely installed in virtual machines and each virtual machine can be turned on and off independently." EX1005, [0005].

62. Hashimoto recognized, however, that "with conventional techniques, no information about routing between virtual machines and virtual storage systems exists, because the server-side virtualization engine and the storage-side virtualization engine do not cooperate with each other." EX1005, [0013]. Accordingly, the object of Hashimoto is "to manage physical paths between a server system and a storage system and information about routing between virtual machines and virtual storage systems in an integrated fashion." EX1005, [0014]. "According to the embodiment of this invention, it is possible to easily realize a high-availability configuration of a computer system including virtual machines and virtual storage systems." EX1005, [0016].

63. Hashimoto describes a system that includes **one or more server systems 100** that can be logically partitioned "to create virtual machines that operate independently" (*see* EX1005, [0052], [0060]), **one or more storage systems 200** that can be logically partitioned "to create virtual storage systems that operate independently" (*see* EX1005, [0052], [0070]), and **control terminal 300** running a virtual machine control program that "controls virtual configuration of the computer system" (*see* EX1005, [0052], [0083]-0084).

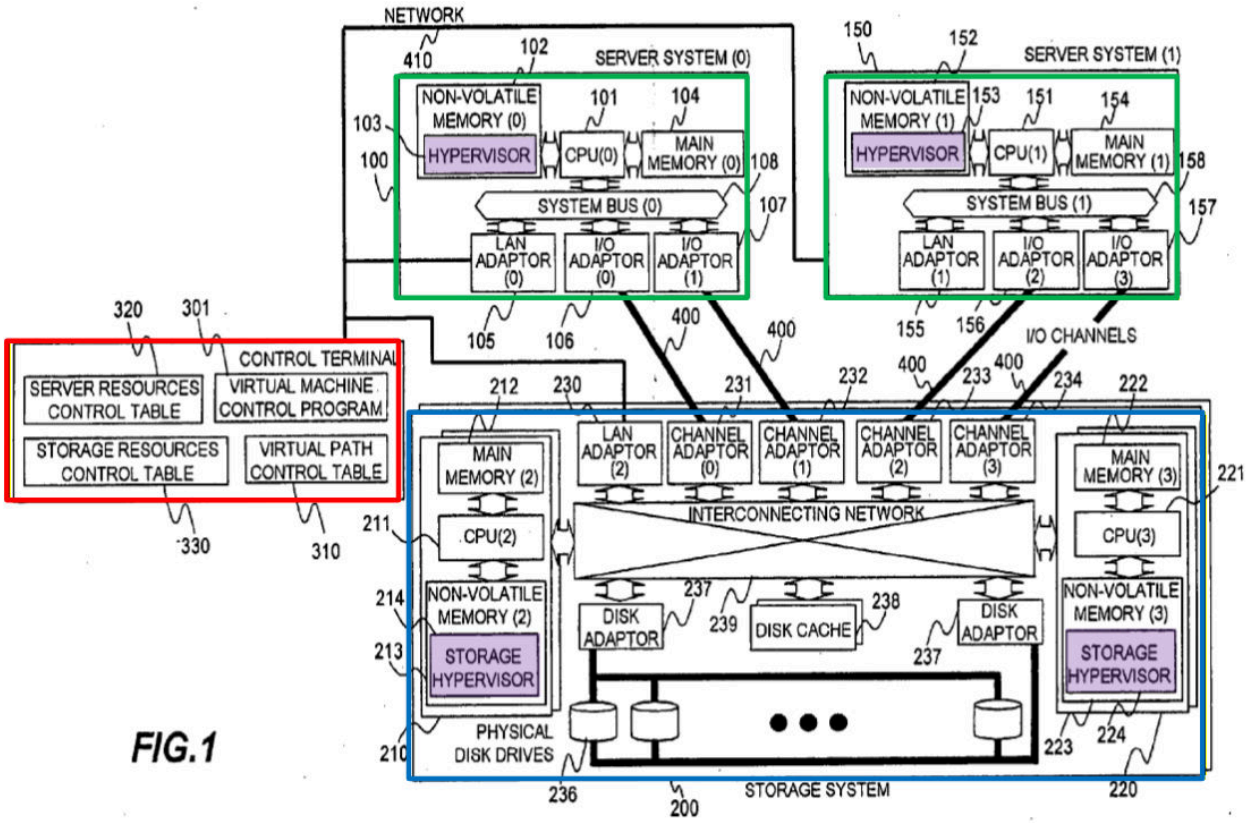


FIG.1

EX1005, Fig. 1 (annotated)

64. The **control terminal 300** is able to control the overall operation of the virtualization system through communication with programs called "**hypervisors**" stored and run on each of the server systems and storage systems. *See* EX1005, [0059]-[0060] ("the hypervisor 103 is a control program that creates logical partitions in the server system (0) 100 to create virtual machines that operate independently"), [0069]-[0070] ("The storage hypervisor 214 is implemented with a control program that creates logical partitions in the storage system 200 to create virtual storage systems that operate independently"). For example, in order to create a virtual machine on a server system 100, an

"administrator operates the control terminal 300 to set allocation of server resources of the server system (0) 100 and gives an instruction to create a virtual machine," and the "server system (0) 100 receives the virtual machine creation instruction from the control terminal 300 and then the hypervisor 103 creates a virtual machine by allocating server resources of the server system (0) 100." EX1005, [0178]-[0179]. Once the virtual machine has been created, the control terminal 300 "checks whether a virtual storage system should be created in the storage system 200," and if so, goes through a similar process with the storage hypervisor to allocate storage resources to a virtual storage system. *See* EX1005, [0181]-[0185].

65. "[I]n the virtual storage system (0) 240 and the virtual storage system (1) 250, the storage hypervisors 214 and 224 may partition a physical disk drive 236 into a plurality of virtual disks 244 and 254, or may unite a plurality of physical disk drives 236 into a single virtual disk 244 or 254." EX1005, [0114]. "The storage system 200 selects a single virtual disk or a plurality of virtual disks from the virtual disks 244 and 254 and provides the virtual disk or disks to the virtual machines 100 and 150 as storage areas." EX1005, [0115]. Application programs run on the virtual machines and the virtual disks store data required for operations of the virtual machines. *See* EX1005, [0052], [0102].

66. Hashimoto describes that, "[w]hen the virtual storage system (0) 240

receives a disk I/O command from the virtual machine (0) 130, the virtual storage system (0) 240 obtains the virtual disk number that is the target of the access by the disk I/O command to specify the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0230]. "Then, the storage hypervisor 214 accesses the physical disk drive 236 and reads/writes data." EX1005, [0230].

67. However, Hashimoto does not limit the manner in which the physical storage is partitioned into virtual storage or the manner in which virtual addresses are translated into physical addresses. *See* EX1005, [0113]-[0116]; [0228]-[0230]. Rather, Hashimoto is more focused on managing the pathways and communication between the virtual machines and virtual storage systems. *See* EX1005, [0014] ("an object of this invention [is] to manage physical paths between a server system and a storage system and information about routing between virtual machines and virtual storage systems in an integrated fashion"), [0052], [0084] ("the control terminal 300 includes a control unit (not shown) having connection information (the virtual path control table 310) that defines connection paths between virtual machines and virtual storage systems"), [0116], [0132]-[0134].

68. Hashimoto is analogous prior art to the '667 patent because it is

directed to the same field of endeavor—virtualization systems—and is reasonably pertinent to the problems the '667 Patent identifies, including the need for more efficient mechanisms for managing virtual environments across computer networks. EX1001, Title, Abstract; EX1005, Title, Abstract, [0014]-[0016].

69. Hashimoto seeks to improve administrators' ability to control, coordinate, and manage virtual machines and associated storage resources, rendering Hashimoto well within the scope of analogous art. EX1001, 3:47-49 ("It would be advantageous if more efficient means for managing virtual environments across computer networks were available"); *See* EX1005, [0014] ("an object of this invention to manage physical paths between a server system and a storage system and information about routing between virtual machines and virtual storage systems in an integrated fashion").

B. Lubbers (EX1006)

70. Lubbers relates to "software, systems and methods for providing and managing virtualized data storage." EX1006, [0002]. "Storage virtualization generally refers to systems that provide transparent abstraction of storage at the block level." EX1006, [0014]. "In essence, virtualization separates out logical data access from physical data access, allowing users to create virtual disks from pools of storage that are allocated to network-coupled hosts as logical storage when needed." EX1006, [0014]. Lubbers recognized that there was a need for "a

virtualization system that operates at a system level," and "a storage allocation system that allocates storage in a manner that provides efficient data structures for accessing the data." EX1006, [0016], [0018].

71. In Lubbers's system, "[p]hysical storage is carved into units called physical segments," called PSEGs. EX1006, [0019]. "Logical storage is implemented in atomic logical units called RStores comprising a range of virtual address space that when allocated, is bound to a particular group of PSEGs." EX1006, [0019]. "A plurality of RStores make up a logical disk that is presented to a user." EX1006, [0019]. Lubber's system includes "a storage pool 101 comprising an arbitrarily large quantity of storage space from which logical disks (also called logical units or LUNs) 102 are allocated." EX1006, [0034]. "Each LUN 102 is essential[ly] a contiguous range of logical addresses that can be addressed by host devices 105, 106, 107 and 109 by mapping requests from the connection protocol used by the hosts to the uniquely identified LUN 102." EX1006, [0035]. "Storage access requests expressed in terms of logical disk addresses are mapped to PSEGs [physical segments of physical storage] containing data represented by the logical addresses through a split-directory representation of the logical unit." EX1006, [0019].

72. The memory mapping data structures proposed by Lubber's are illustrated in Fig. 5 (reproduced below). EX1006, [0024], [0054].

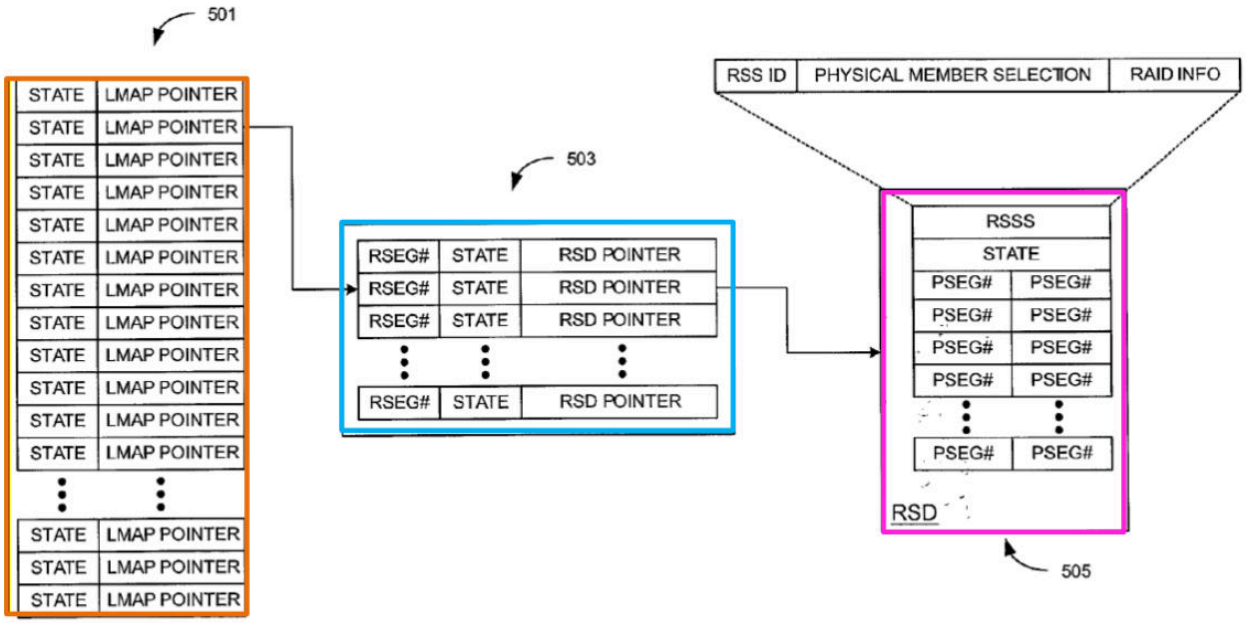


FIG. 5

EX1006, Fig. 5

73. The mapping data structures illustrated in Fig. 5 include **level 2 maps (L2MAP) 501**, **LMAPs 503**, and **redundancy set descriptors (RSDs) 505**. EX1006, [0056]-[0057]. "A single L2MAP handles the entire LUN 102." EX1006, [0057]. "L2MAP 501 includes a plurality of entries where each entry represents 2 Gbyte of address space." EX1006, [0058]. "An LMAP 503 exists only for each 2 Gbyte increment of allocated address space." EX1006, [0060]. In other words, "[e]ach LUN 102 is represented by multiple LMAPs 503 where the particular number of LMAPs 503 depend on the actual address space that is allocated at any given time." EX1006, [0057].

74. "An LMAP 503 exists only for each 2 Gbyte increment of allocated

address space." EX1006, [0060]. "Each LMAP 503 includes a plurality of entries where each entry optionally corresponds to a redundancy segment (RSEG)." EX1006, [0060]. "An RSEG is an atomic logical unit that is roughly analogous to a PSEG in the physical domain—akin to a logical disk partition of an RStore." EX1006, [0060]. "Each LMAP entry also includes state information about the particular RSEG, and an RSD pointer." EX1006, [0060]. "The RSD pointer points to a specific RSD 505 that contains metadata describing the RStore in which the corresponding RSEG exists." EX1006, [0062].

75. "An RStore is analogous to a RAID disk set, differing in that it comprises PSEGs rather than physical disks." EX1006, [0068]. "An RStore is smaller than conventional RAID storage volumes, and so a given LUN 102 will comprise multiple RStores as opposed to a single RAID storage volume in conventional systems." EX1006, [0068]. A primary benefit of the RStore/RSEG logical structure is that Lubbers's system "can readily form a nested RAID set using, for example, increments of RStores rather than increments of six drives." EX1006, [0075]. "Because RStores are much smaller than typical drives, the granularity with which capacity can be added and removed from a LUN 102 is greatly improved." EX1006, [0075]. "A LUN 102 can be incremented by adding an RStore and consume as little as 8 Mbyte of additional capacity." EX1006, [0075].

76. Lubbers is analogous prior art to the '667 patent because it is directed to the same field of endeavor—virtualization systems—and is reasonably pertinent to the problems the '667 Patent identifies, such as providing more efficient mechanisms for managing virtual environments across computer networks. EX1001, Title, Abstract; EX1006, Title, Abstract, [0012]-[0014].

77. Both the '667 Patent and Lubbers seek to improve administrators' ability to manage virtualized resources, with Lubbers expressly describing virtualization implemented in a storage system controller to deliver high performance, high data availability, fault tolerance, and efficient storage management. Accordingly, Lubbers falls squarely within the scope of analogous art. EX1001, 3:47-49 ("It would be advantageous if more efficient means for managing virtual environments across computer networks were available"); EX1006, [0030] ("virtualization in accordance with the present invention is implemented in a storage system controller to provide high performance, high data availability, fault tolerance, and efficient storage management").

C. Hirschfeld (EX1007)

78. Hirschfeld is directed to a "virtualized logical server cloud" in which users access logical servers through a portal-like user interface without needing to know the access information of the underlying physical servers. EX1007, [0008], [0031], [0066]. Although virtualization allows "multiple logical servers to

override a physical server's attributes and share its resources" (*id.*, [0005]) users historically "had to be aware of and have specific information about the underlying physical server to access the logical servers" (*id.*, [0006]), such as the physical server's IP address/server name and port number, and incorrect access inputs could result in access being denied. EX1007, [0018], [0031], [0066].

79. Hirschfeld identifies a desire "to completely separate logical servers from physical servers so that there is no permanent tie between physical and logical resources," enabling physical servers to function as a resource pool that supports logical servers and allowing a logical server to be reallocated to a different physical server "without users experiencing any change in access approach." EX1007, [0008].

80. Consistent with this objective, Hirschfeld describes a server cloud manager ("SCM") that provides a "web page portal or other access mechanism" for remote terminals to communicate with the logical server cloud, such that "a user of any remote terminal" may "login to and access any of the logical servers" via the user interface and SCM. EX1007, [0030]. A user may employ a "local browser to access a web page" associated with the user interface, and the GUI provides a "portal or other interface" through which the user may access one or more logical servers based on permission (e.g., by "point and click" selection of a logical server graphic). *Id*., [0031].

81. Hirschfeld also describes a user database that "incorporates a mapping of users, the logical servers and the physical servers," including "link information that maps each logical server" with a corresponding physical server, such that the SCM "automatically accesses the appropriate physical server corresponding to the selected logical server," and the user "does not have to supply physical server information" and "does not even need to be aware of" physical-server access/location information. EX1007, [0032]-[0033]. Hirschfeld emphasizes that its virtualization "enables a logical server to exist independent of a specific physical server," with logical attributes "non-deterministically allocated to physical resources," thereby creating a cloud of logical servers over one or more physical servers. *Id*., [0065].

82. Hirschfeld is analogous prior art to the '667 patent because it is directed to the same field of endeavor—virtualization systems—and is reasonably pertinent to the problems the '667 Patent identifies, such as providing more efficient mechanisms for managing virtual environments across computer networks. EX1001, Title, Abstract; EX1007, Title, Abstract, [0008].

83. Both the '667 Patent and Hirschfeld address improving users' and administrators' ability to efficiently manage a virtualization environment. Like the '667 Patent, Hirschfeld describes virtualization implemented within a storage system controller to enhance logical server management, delivering high

performance, high data availability, fault tolerance, and efficient storage management. *See* EX1001, 3:47–49 ("It would be advantageous if more efficient means for managing virtual environments across computer networks were available"); EX1007, [0067] ("Logical server management is substantially improved employing a virtualized logical server cloud in accordance with embodiments of the present invention"). Hirschfeld therefore falls squarely within the scope of analogous art.

D. Khandekar (EX1012)

84. Like Hashimoto, Khandekar teaches "a method and system…for creating a virtualized computer system based on input information identifying a desired computer configuration." EX1012, 4:30-35; *see* EX1005, FIG. 10, [0177]-[0180] (describing creation of a virtual machine). "Of course, most computers are intended to run various applications, and VMs are usually no exception." EX1012, 6:24-27. Khandekar teaches, as a part of its invention, that "the requester specifies various components of the desired VM, including the characteristics of a desired operating system (OS), of a desired hardware platform, of desired applications, or of any combination of these components." EX1012, 4:42-46. And, "[o]nce a VM has been configured and deployed," Khandekar's invention makes possible to "alter its configuration, for example by upgrading the software and/or hardware, installing different applications, etc." EX1012, 4:48-51.

85. Specifically, "[t]he deployed VMs can be monitored regularly so that any unexpected event, like the crash of the operating system in the VM, can be reported to the user or administrator." EX1012, 13:44-46. "Also, maintenance tasks like upgrade of software and hardware (for example, activating a different number of processors, memory chips, etc.) installation of service packs or hot fixes, detection of virus, etc., can be run by the administrative component of the provisioning server." EX1012, 13:46-51. "Hosts can be monitored for load and if the load crosses a threshold, the administrator can be alerted and/or actions can be taken such as migrating VMS to a lightly loaded server." EX1012, 13:51-54.

86. These tasks of monitoring deployed VMs is facilitated by virtual machine monitors (VMMs). *See* EX1012, 6:62-8:44 (describing VMM functions). "A VMM is usually a thin piece of software that runs directly on top of the host and virtualizes all, or at least some of, the resources of the machine, or at least of some machine." EX1012, 7:7-9. In other words, Khandekar's VMM performs a similar role to Hashimoto's hypervisor. *Compare id. with* EX1005, [0060].

87. Khandekar is analogous prior art to the '667 patent because it is directed to the same field of endeavor—virtualization systems—and is reasonably pertinent to the problems the '667 Patent identifies, such as providing more efficient mechanisms for managing virtual environments across computer networks. EX1001, Title, Abstract; EX1012, Title, Abstract, 1:7-10.

88. Both the ’667 Patent and Khandekar address improving users’ and administrators’ ability to efficiently manage a virtualization environment. Like the ’667 Patent, Khandekar describes a virtualization environment that allows for selection of preferred virtual system configuration and recovery from virtual system failures. *See* EX1001, 3:47–49 (“It would be advantageous if more efficient means for managing virtual environments across computer networks were available”), 9:12-15 (“the Optimizer component may operate on a user desktop or on a server that manages a virtual environment so as to configure an efficient combination of applications and assets, such as virtual devices of the host machine”), 12:28-30 (“If a particular VM fails, the Control Agent 324 tries to restart the VM, or if that fails, sends out an alarm message to the virtual server”); EX1012, 22:14-33 (“a method and system implementation for creating a virtualized computer system based on input information identifying a desired computer configuration”), 13:44-46 (“[t]he deployed VMs can be monitored regularly so that any unexpected event, like the crash of the operating system in the VM, can be reported to the user or administrator”)

X. OVERVIEW OF THE PRIOR ART COMBINATIONS

A. The Combination of Hashimoto and Lubbers

89. In my opinion, Hashimoto does not prescribe any particular technique for partitioning physical storage into virtual storage, nor does it require a specific

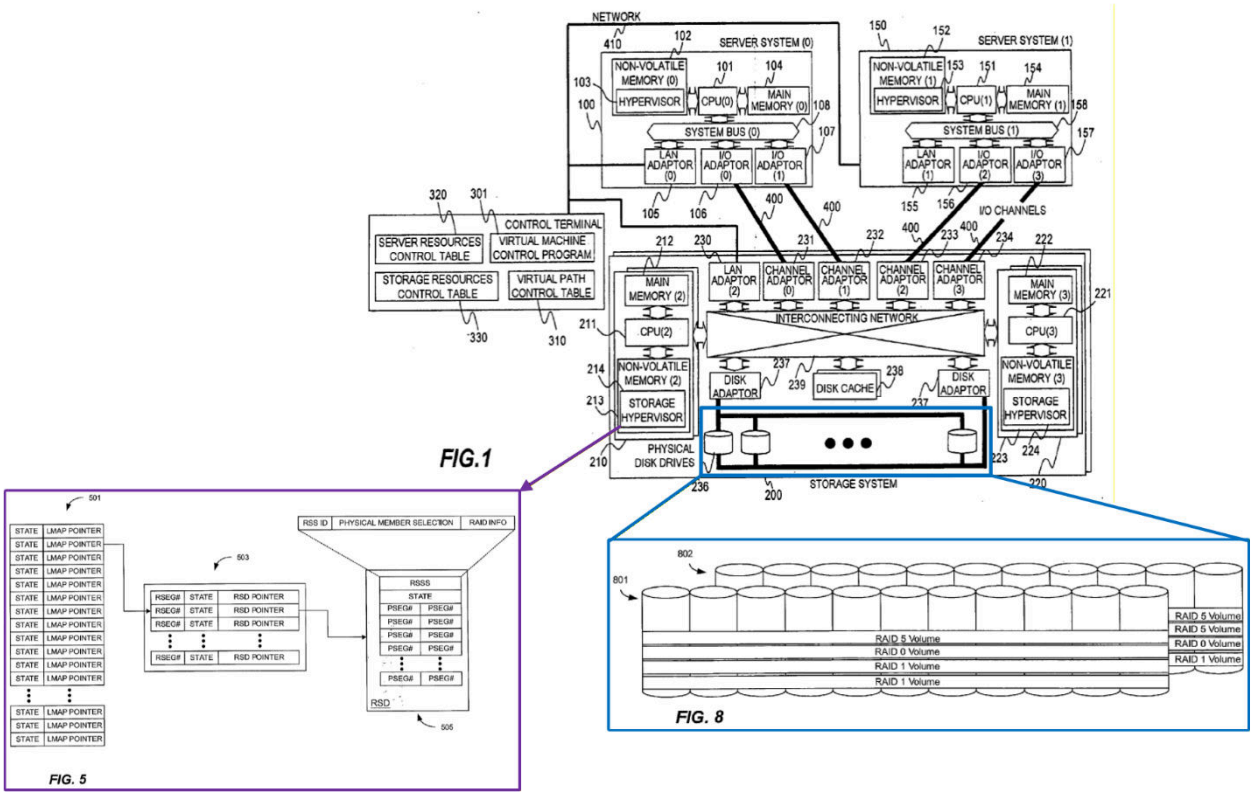
mechanism for translating virtual addresses to physical addresses. *See* EX1005, [0113]-[0116]; [0228]-[0230]. Rather, Hashimoto focuses on managing the pathways and communications between virtual machines and virtual storage systems. *See* EX1005, [0014], [0052], [0084], [0116], [0132]-[0134]. Consistent with that focus, Hashimoto leaves the implementation details of how to partition physical storage into virtual storage to a POSITA implementing the system.

90. Lubbers provides a useful, concrete example of how to carry out such partitioning within Hashimoto's virtualization environment. A POSITA would have found it obvious to incorporate the partitioning techniques described by Lubbers into Hashimoto's system. For example, Lubbers discloses a SAN manager appliance 109 that manages a "quorum space 113" in which metadata about the overall storage system is maintained. *See* EX1006, [0040]-[0041]. A POSITA would have recognized that Lubbers's SAN manager appliance 109 is similar in structure and purpose to Hashimoto's control terminal 300, which likewise maintains data tables used to manage the virtualization environment, including a storage resources control table 330 that "stores the associations of the resources of the storage system 200." *See* EX1005, [0083]-[0085], [0116], [0161]. As another example, Lubbers discloses storage cells 203, each with a redundant pair of NSCs 301 that "implement[] a network interface to receive storage access requests from hosts as well as fibre channel arbitrated loop ports to connect to

storage device in cabinet 303." *See* EX1006, [0049]-[0051]. A POSITA would have recognized that Lubbers's storage cell 203 and NSCs 301 are similar in structure and function to Hashimoto's storage system 200 and redundant control blocks 210 and 220, which "control data input/output to and from physical disk drives 236 and manages operations of the storage system 200." *See* EX1005, [0065]-[0066]. Indeed, Hashimoto expressly teaches that "one virtual storage system can be configured by extracting resources of a plurality of storage systems," such as the collection of storage cells 203 in Lubbers, and that its invention is applicable "to such… virtual storage system configurations." EX1005, [0374]-[0375]. In short, Hashimoto already provides much of the underlying structure that would accommodate Lubbers's partitioning techniques.

91. Under the proposed combination (a diagram of which is included below), Hashimoto's virtualization system would implement the various data structures and control functionality described by Lubbers to virtualize Hashimoto's storage system 200 and to "enable[] translation of a request expressed in terms of a logical block address (LBA) from host 301 into a read/write command addressed to a particular portion of a physical disk 405 having the desired information," as described in Lubbers and consistent with Hashimoto's framework. *See* EX1006, [0030]-[0033], [0046], [0054], [0127]-[0128]. In the example implementation shown below, Hashimoto's physical disk drives 236 would have stored the **storage**

volumes described by Lubbers, and Hashimoto's storage hypervisors would store the **mapping structures of the LUNs** described in FIG. 5 of Lubbers. *See id.* A POSITA would have been motivated to combine Hashimoto and Lubbers for at least the following reasons.



Amalgamation of Fig. 1 of EX1005 and Figs. 5 and 8 of EX1006

92. First, Lubbers explains that its virtualization method "atomizes storage capacity in the physical domain in the form of PSEGs and in the logical domain in the form of RStores to enable fluid movement of capacity," and that "[t]he relationship of particular disks can be modified to provide various levels of protection and allow for construction of virtual disks with variable levels of data

protection." EX1006, [0128]. A POSITA would have recognized that these express benefits are directly applicable to Hashimoto's similarly structured storage system 200.

93. Second, integrating Lubbers's storage-virtualization approach into Hashimoto's system would have advanced Hashimoto's stated goals. Hashimoto aims to "realize a high-availability configuration of a computer system including virtual machines and virtual storage systems." EX1005, [0016]. Lubbers likewise teaches that its method "provide[s] high performance, high data availability, fault tolerance, and efficient storage management." EX1006, [0030].

94. Third, Lubbers provides a known, concrete example of storage virtualization techniques not otherwise specified in Hashimoto. As such, the combination reflects the obvious application of a known technique (Lubbers's virtualization of networked storage) to similar devices (Lubbers's storage cells 203, each with a set of physical disks, and Hashimoto's storage system 200, which also includes a set of physical disks) to achieve the same type of improvement.

95. A POSITA also would have reasonably expected success implementing this combination. The resulting system would be realized using conventional software and hardware (e.g., general-purpose processors and storage systems interconnected over well-known networks) and well-understood virtualization and data-storage techniques (e.g., RAID) that were widely known by

the Critical Date. Further, the techniques integrated from Lubbers in the Hashimoto–Lubbers combination are fully compatible with Hashimoto's computer system of FIG. 1 and would not disturb the ability of Horvath's system to transmit or deliver SMS messages over either a packet-based or a circuit-switched network.

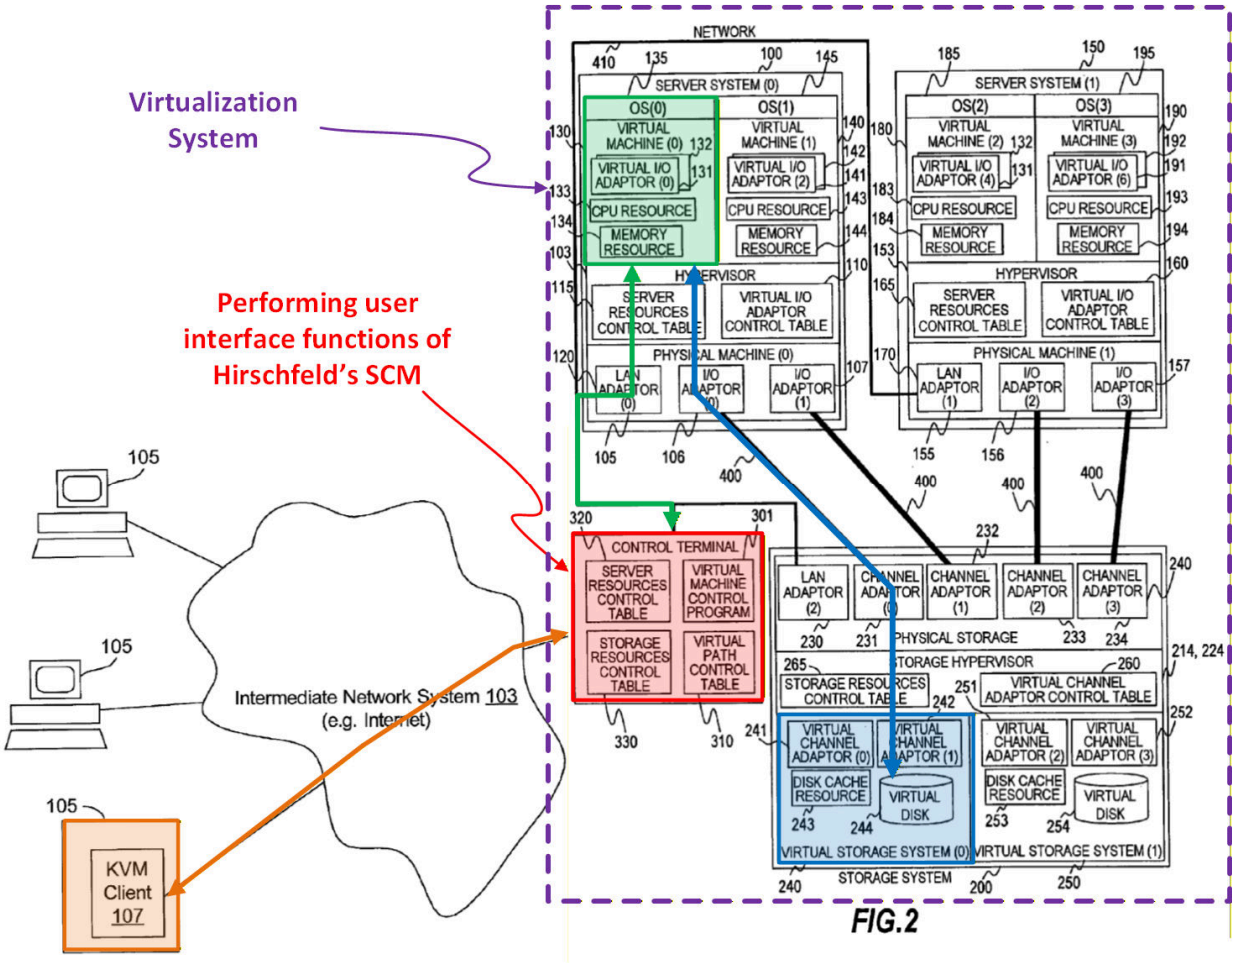
B. The Combination of Hashimoto, Lubbers, and Hirschfeld

96. In my opinion, Hashimoto—like Hirschfeld—describes partitioning physical resources into virtual machines. *See* EX1005, [0059]-[0060], [0069]-[0070]; EX1007, [0028], [0039]. And, similar to Hirschfeld's SCM, Hashimoto discloses a distinct control terminal (element 300) that governs the overall virtualization system. *See* EX1005, [0052] ("The control terminal 300 manages information for controlling the computer system and controls operations of the entire computer system."); EX1007, [0040] ("the SCM 102 is implemented on a separate physical server or administrator computer," and "includes SCM software to manage the logical servers 106 and the physical servers 108"). Unlike Hirschfeld, however, Hashimoto does not explicitly describe a mechanism by which an end user would interface with the system to utilize virtual machines and virtual storage. Rather, Hashimoto focuses on how an administrator interfaces with the control terminal 300 to manage the system. *See* EX1005, [0310]-[0311] ("The administrator can manage the computer system by operating the control terminal 300."), [0321]-[0322].

97. Hirschfeld provides an obvious solution for how users would interface with the combined Hashimoto–Lubbers virtualization system. Hirschfeld expressly identifies that, in conventional virtualization deployments, users "had to be aware of and have specific information about the underlying physical server" (e.g., IP address/server name and port number) to access logical servers, which created usability and operational issues (including the need for distinct access information for each logical server and updated information when servers are moved). EX1007, [0008]. Hirschfeld further explains that it is "desired to completely separate logical servers from physical servers" so that logical servers may be reallocated to different physical servers "without users experiencing any change in access approach." *Id*., [0008].

98. In addressing these user-access and mapping challenges in a virtualized environment, a POSITA would have found Hirschfeld's disclosure of a "web page portal or other access mechanism" instructive. Through an SCM and user interface, Hirschfeld enables remote terminals to log in to and access logical servers in a logical-server cloud via a browser-based portal. EX1007, [0030]-[0031]. Hirschfeld further describes that the SCM maintains a user database incorporating a "mapping of users, the logical servers and the physical servers," including link information mapping each logical server to a physical server, such that the SCM "automatically accesses the appropriate physical server

corresponding to the selected logical server," eliminating any need for the user to supply (or even be aware of) physical-server access information. *Id*., [0032].



Amalgamation of EX1005 Fig. 2 and EX1007 Fig. 1

99. A POSITA would have viewed Hirschfeld's SCM/portal architecture as a compatible enhancement to the Hashimoto–Lubbers control console (Hashimoto's control terminal 300), because both provide network-accessible control-plane interfaces for interacting with virtualized server instances. *Compare* EX1005, [0052] ("The control terminal 300 manages information for controlling

the computer system and controls operations of the entire computer system.") *with* EX1007, [0040] ("the SCM 102 is implemented on a separate physical server or administrator computer," and "includes SCM software to manage the logical servers 106 and the physical servers 108").

100. Accordingly, a POSITA would have been motivated to leverage Hirschfeld's SCM-managed mapping and portal-based access mechanisms within the Hashimoto–Lubbers virtualization system to provide a user-facing experience in which end users interact through a browser-based console/portal while the system maintains and uses the mapping between logical server instances (e.g., provisioned VMs/logical servers) and the underlying physical servers. *See* EX1007, [0031]-[0033]. By implementing the user interface functions of Hirschfeld's SCM with Hashimoto's control terminal (a "SCM/control terminal"), Hashimoto-Lubbers-Hirschfeld advances Hirschfeld's stated objective of removing any requirement that users know physical-server-specific access information, while remaining consistent with Hashimoto-Lubbers's architecture in which the control terminal 300 provides an interface to "manage the computer system," and is connected to the server systems 100 and 150 and storage system 200 through a network 410. *See* EX1005, [0083]-[0087], [0310]-[0311], [0321]-[0322]; *see also* EX1006, [0030]-[0033], [0066]. A benefit of this combination would have been a simpler and more efficient user experience with the virtualization system.

101. A POSITA would have reasonably expected success in combining the teachings of Hashimoto, Lubbers, and Hirschfeld in the manner discussed above. Each reference addresses virtualized computing environments in which centralized control-plane components coordinate the provisioning and/or access to logical server instances, and each describes interfaces for system control. *See* EX1005, [0083]-[0087], [0310]-[0311], [0321]-[0322]; EX1006, [0040] (describing the SAN manager appliance 109); EX1007, [0030]-[0033]. The proposed combination applies Hirschfeld's SCM portal-access framework to the Hashimoto–Lubbers control terminal— which already serves a similar system-control function to Hirschfeld's SCM. *See* EX1007, [0029]-[0030]; EX1005, [0052], [0083]-[0087]. In my opinion, a POSITA would have understood this as combining known elements according to known methods to yield predictable results.

C. The Combination of Hashimoto, Lubbers, and Khandekar

102. Hasimoto describes that "[a]pplication programs run in the server system (0) 100 and the server system (1) 150." EX1005, [0052]. These "[a]pplication programs run on the OSs 135 and 145" of the virtual machines. EX1005, [0102]. "These application programs provide database services, Web services, etc. to clients (not shown) connected to the server system (0) 100." EX1005, [0102]. However, Hashimoto does not describe how the applications are distributed to the virtual machines or otherwise managed within the virtualization

environment. Khandekar provides details about at least one obvious way in which Hashimoto's virtualization system would distribute and manage the application programs run on the OSs of the virtual machines.

103. Based on the teachings of Khandekar I note above in §IX.D, a POSITA would have found it obvious to apply the application selection, deployment, and management techniques described by Khandekar to Hashimoto's control terminal and hypervisors. For example, a POSITA would have found it obvious to utilize Hashimoto's hypervisor and control terminal to perform the application deployment and management functions described by Khandekar with respect to its VMM and provisioning server. *See* EX1012, 7:7-27, 8:20-36, 22:14-33. A POSITA would have been motivated to combine Khandekar with Hasimoto and Lubbers in this manner, because the deployment and management functions described by Khandekar would have created greater flexibility within the Hasimoto-Lubbers system (*e.g.*, allowing users to freely select applications for use on the virtual machines and change them after deployment), increased resilience (*e.g.*, detecting and recovering from faults and upgrading older software), and increased performance (*e.g.*, real-time load balancing). *See* EX1012, 4:42-51, 13:44-54.

XI. THE PRIOR ART MAKES CLAIMS 1-20 OBVIOUS

A. Ground 1: Claims 1-5, 7, and 8 Are Obvious Over The Combination of Hashimoto and Lubbers

1. *Claim 1*

[1pre]

104. Hashimoto describes a system and associated methods that manage the physical paths between server systems and storage systems, along with information about routing between virtual machines and virtual storage systems, all in an integrated manner. EX1005, [0014]. In particular, Hashimoto discloses a computer system composed of server systems 100 and 150, at least one storage system 200, and a control terminal 300, all interconnected through a network 410. EX1005, [0052]; *see also* [0371]-[0375]. The storage system 200 stores the data used by the server systems 100 and 150. EX1006, [0052]. A program referred to as a "storage hypervisor" runs on the storage system 200, and this program creates logical partitions within the storage system so that independent virtual storage systems can be formed. EX1005, [0070].

105. As explained above with respect to the combination of Hashimoto and Lubbers, when Hashimoto is combined with Lubbers, Hashimoto's storage virtualization architecture would rely on Lubbers's data structures and control mechanisms for virtualizing the storage system 200. Lubbers's disclosure, in turn, relates to software, systems, and methods for providing and managing virtualized data storage. EX1006, [0002].

[1.1.a]

106. As described in connection with [1.pre], Hashimoto's computer system includes servers 100 and 150, at least one storage system 200, and a control terminal 300, interconnected through network 410. EX1005, [0052]; *see also* [0371]-[0375]. A virtual machine control program 301 executes on the CPU of the control terminal 300. EX1005, [0083]. That control program manages the virtual configuration of the system using a virtual path control table (310), a server resources control table (320), and a storage resources control table (330). EX1005, [0084].

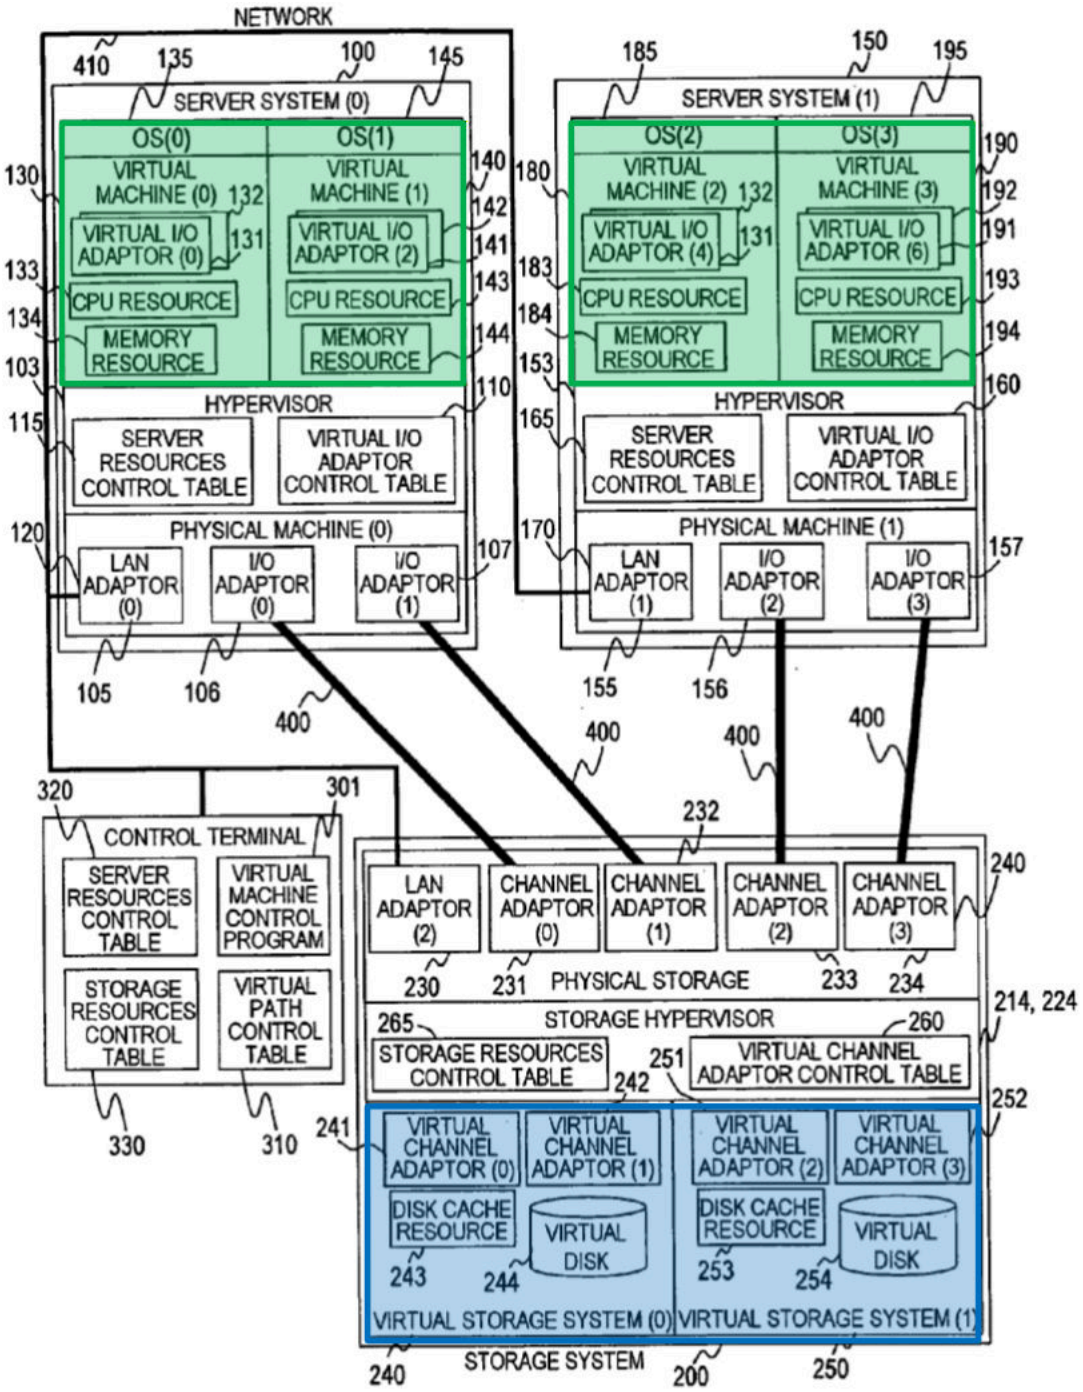


FIG.2

EX1005, FIG. 2 (Annotated)

107. Each server system is "a computer having physical resources including a CPU (0) 101, a non-volatile memory (0) 102, a main memory (0) 104,

a LAN adaptor (0) 105, an I/O adaptor (0) 106, an I/O adaptor (1) 107, and a system bus 108." EX1006, [0053]. "The non-volatile memory (0) 102 stores a hypervisor 103." EX1006, [0059]. "The hypervisor 103, implemented by the execution of processing by the CPU (0) 101, is a means that enables logical partitioning of physical resources of the server system (0) 100." EX1006, [0059].[1] The "hypervisor 103 is a control program that creates logical partitions in the server system (0) 100 to create **virtual machines** that operate independently." EX1006, [0060].

108. Likewise, "[t]he storage system 200 has a storage control block 210 including a CPU (2) 211, a main memory (2) 212, and a non-volatile memory (2) 213." EX1005, [0065]. A second storage control block 220 includes similar components. EX1005, [0065]. "The storage control blocks 210 and 220 control data input/output to and from physical disk drives 236 and manages operations of the storage system 200." EX1005, [0066]. Storage hypervisor 214 and 224 are "implemented by the execution of processing by the CPUs 211 and 221, is a means that enables logical partitioning of physical resources of the storage system 200." EX1005, [0069]. "The storage hypervisor 214 is implemented with a control program that creates logical partitions in the storage system 200 to create **virtual**

[1] Emphasis added as underline throughout, unless noted otherwise.

storage systems that operate independently." EX1005, [0070].

109. The control terminal 300 manages information for controlling the computer system and controls operations of the entire computer system." EX1005, [0052]. "A virtual machine control program 301 runs in the CPU of the control terminal 300." EX1005, [0083]. "The virtual machine control program 301 controls virtual configuration of the computer system by using the virtual path control table 310, the server resources control table 320, and the storage resources control table 330." EX1005, [0084]. The control program 301 communicates with the hypervisors 103 and 153 and storage hypervisors 214 and 224 to control the partitioning and maintain the control tables. *See* EX1005, [0060], [0070], [0179]-[0180], [0184], [0193].

110. FIG. 2 (above) shows each of server systems 100 and 150 and storage system 200 divided into a physical layer, a hypervisor layer, and a virtual machine/storage layer. EX1005, [0090], [0104]. Server systems 100 and 150 have each been partitioned into two virtual machines (130, 140, 180, and 190). EX1005, [0100], [0103], [0164]-[0167]. Similarly, storage system 200 has been partitioned into two virtual storage systems (240 and 250). EX1005, [0113]-[0114].

111. Accordingly, Hashimoto-Lubbers ***implement[s] a virtualization system*** (*e.g.*, Hashimoto's computer system of FIGS. 1 and 2) ***that includes a***

control center application (*e.g.*, virtual machine control program 301)***, one or more physical resources*** (*e.g.*, the physical resources of server systems 100 and 150, and storage system 200)***, and one or more virtual assets on each of the physical resources*** (*e.g.*, virtual machines 130, 140, 180, and 190 and virtual storage systems 240 and 250).[2]

[1.1.b]

112. As discussed for [1.1.a], *supra*, the server systems 100 and 150 run a hypervisor 103, 153, and the storage control blocks 210 and 220 of storage system 200 run a storage hypervisor 214, 224. *See* EX1005, Fig. 1-2, [0059]-[0060], [0069]-[0070], [0092], [0103], [0106]. As shown in FIG. 2, "[f]unctionally, the server system (0) 100 can be roughly divided into a physical layer, a hypervisor layer, and a virtual machine layer." EX1005, [0090]. The "hypervisor layer is implemented by the hypervisor 103…" EX1005, [0092]. Similarly, "storage system 200 can be roughly divided into a physical layer, a hypervisor layer, and a virtual storage layer." EX1005, [0104]. "The hypervisor layer is implemented by the storage hypervisors 214 and 224…" EX1005, [0106].

[2] Claim language is shown in bold and italics.

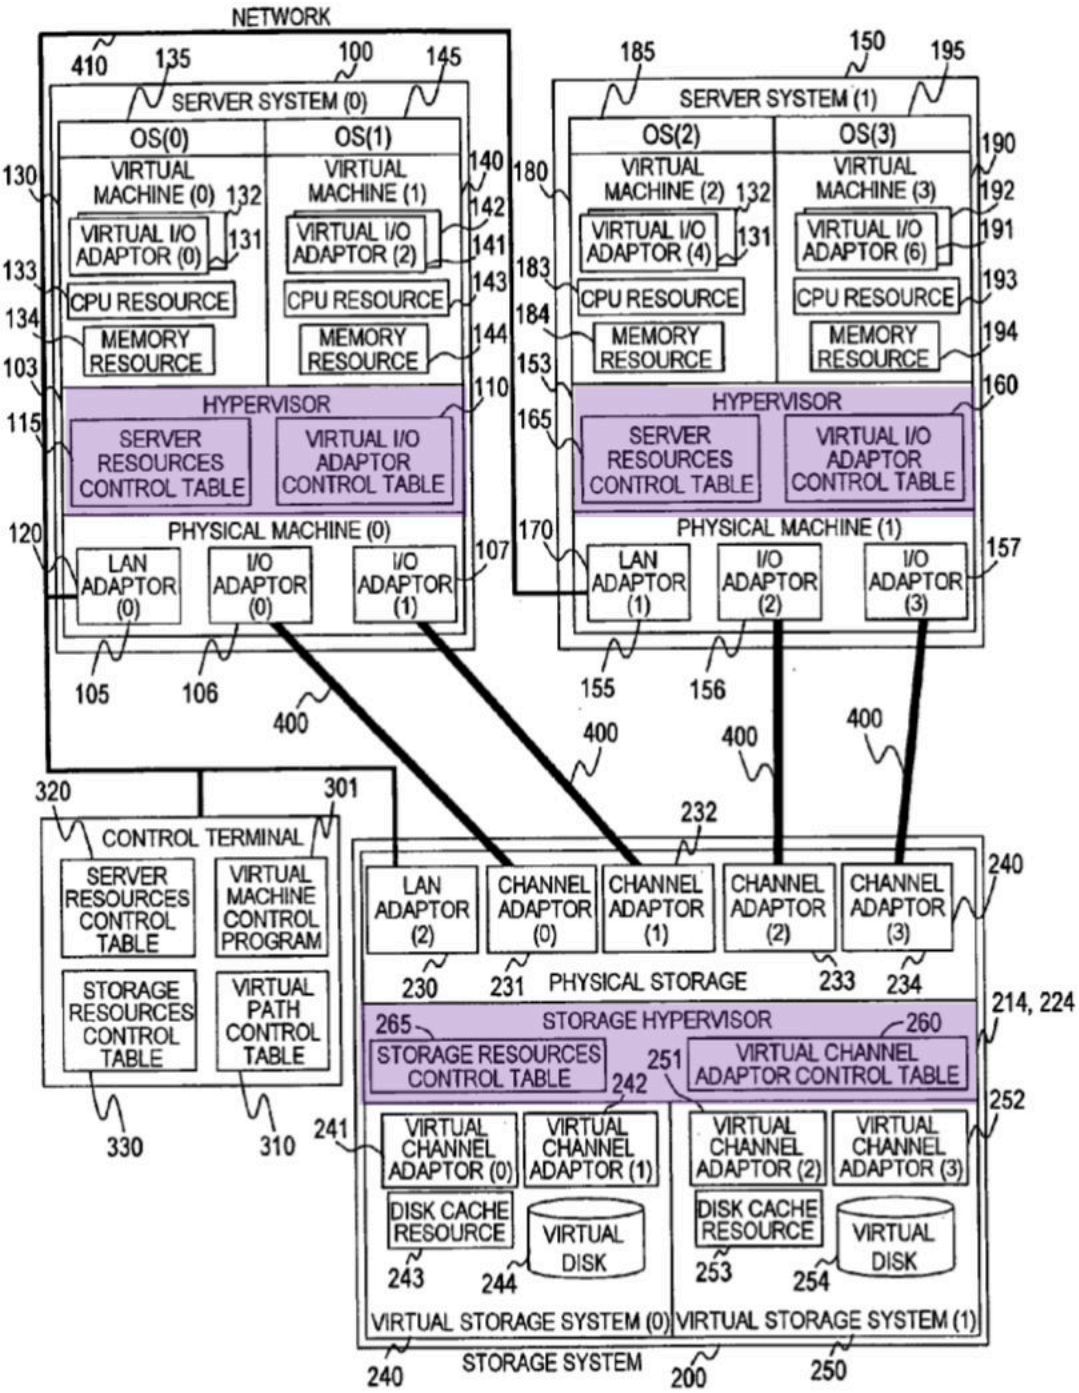


FIG.2

EX1005, FIG. 2 (Annotated)

113. The hypervisors 103 and 153 perform management functions including, for example, creating the virtual machines (*see* EX1005, [0179]-[0180]),

routes I/O requests generated by the virtual machines (*see* EX1005, [0223]-[0224]), detecting and fixing faults in the I/O channels of the virtual machines (*see* EX1005, [0236]-[0246]), and rebooting the virtual machines (*see* EX1005, [0291]-[0296]). The storage hypervisors 214 and 224 perform similar management functions including, for example, creating the virtual storage systems (*see* EX1005, [0184]-[0185]), routes I/O requests generated by the virtual machines (*see* EX1005, [0206], [0226]-[0227]), and detecting and fixing faults in the I/O channels of the virtual machines (*see* EX1005, [0247]-[0259]).

114. Accordingly, in Hashimoto-Lubbers, ***the virtual assets*** (*e.g.*, virtual machines 130, 140, 180, and 190 and virtual storage systems 240 and 250) ***are managed by virtualization layer instances*** (*e.g.*, the hypervisor layers implemented by hypervisors 103 and 153 and storage hypervisors 214 and 224) ***running on each of the physical resources***.

[1.1.c]

115. As discussed above with respect to [1.1.a], storage system 200 has been partitioned into two virtual storage systems (240 and 250). EX1005, [0113]-[0114].

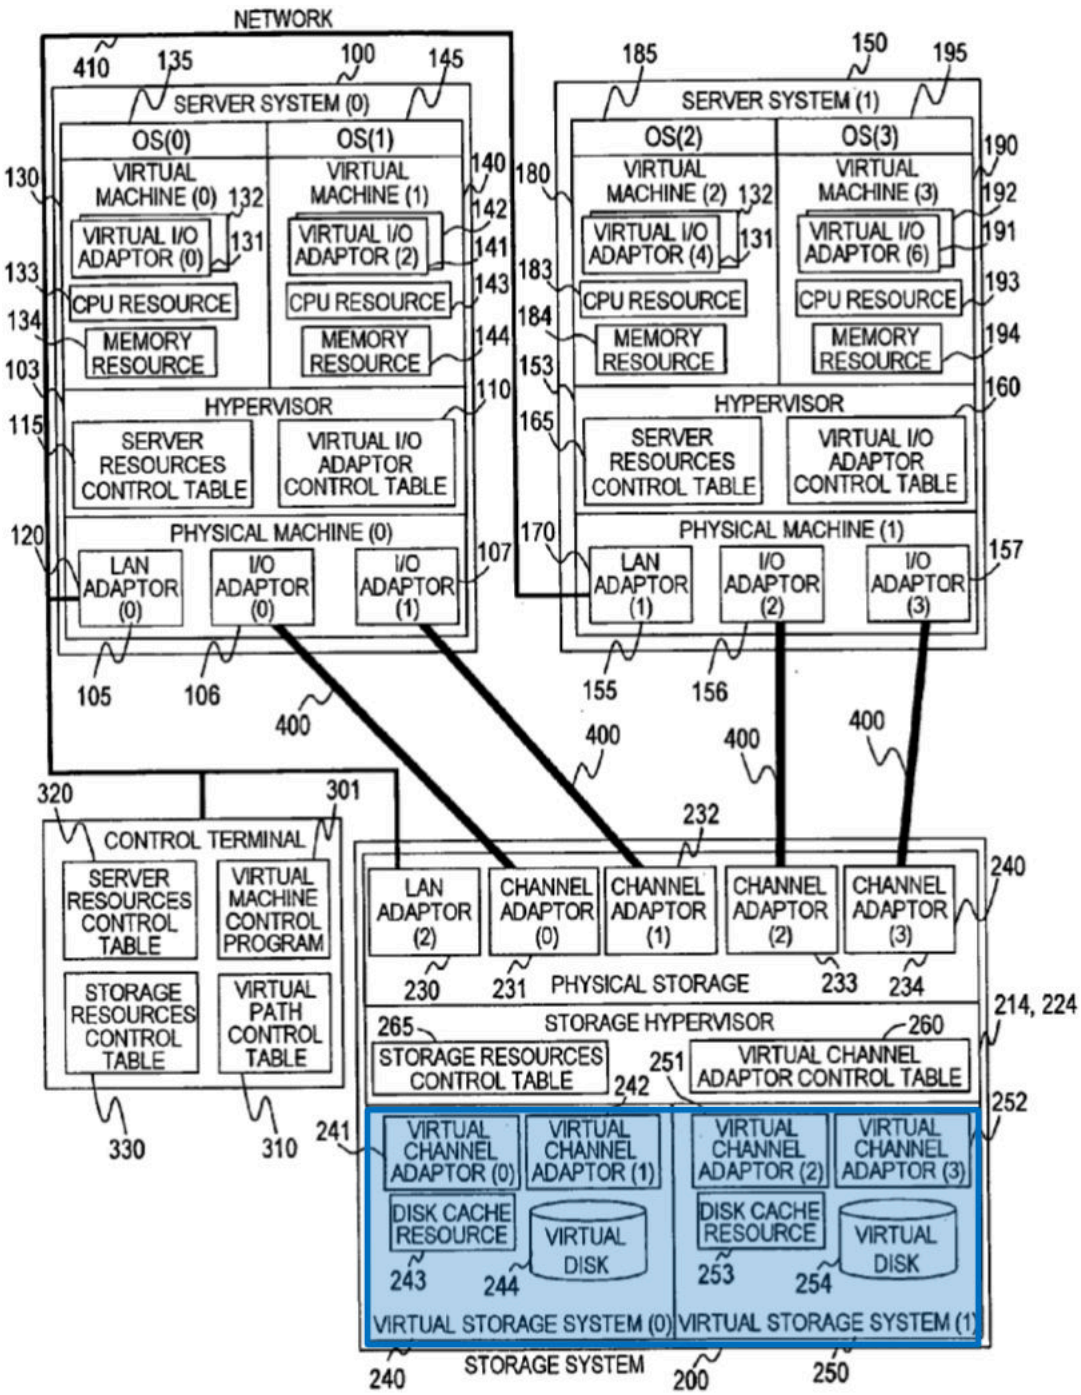


FIG.2

EX1005, FIG. 2 (Annotated)

116. Accordingly, in Hashimoto-Lubbers, ***at least two of the virtual assets on one or more of the physical resources*** (*e.g.*, storage system 200) ***include***

virtual storage assets (*e.g.*, virtual storage systems 240 and 250).

[1.2]

117. FIG. 14 of Hashimoto "is a flowchart of a disk I/O command transmission process," and more particularly "an example in which the virtual machine (0) 130 accesses the virtual storage system (0) 240 through the I/O adaptor (0) 106 and the channel adaptor (0) 231." EX1005, [0222]. "When a disk I/O command is issued, the OS (0) 135 sends the disk I/O command to the virtual machine (0) 130 (the virtual I/O adaptor (0) 131)," which is "received by the hypervisor 103." EX1005, [0222]. The I/O command "specif[ies] the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. The storage system 200 receives the I/O command via channel adaptor (0) 231. EX1005, [0226]-[0227].

118. As part of fulfilling the I/O command, "the storage hypervisor 214 translates the virtual block address of the accessed virtual dis 244 to a physical block address of the physical disk drive 236." EX1005, [0230]. Lubbers similarly describes "translation of a request expressed in terms of a logical block address (LBA) from host 301 into a read/write command addressed to a particular portion of a physical disk 405 having the desired information." EX1006, [0054]. As discussed above in §X.A, Lubbers provides, for example, the translation methods utilized in Hashimoto-Lubbers.

119. Accordingly, Hashimoto-Lubbers ***obtain[s] a set of data to be stored*** (*e.g.*, a virtual storage system obtains data included in an I/O write command that is sent from a virtual machine, and the virtual storage system "writes data to… the corresponding virtual disk") ***the virtual storage assets on any of the physical resources*** (*e.g.*, a virtual disk of a virtual storage system run on storage system 200).

[1.3]

120. As discussed above with respect to the combination of Hashimoto and Lubbers, Hashimoto–Lubbers would implement the various data structures and control functionality described by Lubber to virtualize Hashimoto's storage system 200 and "enable[] translation of a request expressed in terms of a logical block address (LBA) from host 301 into a read/write command addressed to a particular portion of a physical disk 405 having the desired information," as taught by Hashimoto. *See* EX1006, [0030]-[0033], [0046], [0054], [0127]-[0128].

121. Lubbers further describes that logical disks (also called logical units or LUNs) 102 are allocated from a storage pool provided by a set of physical storage devices. EX1006, [0034]-[0035]. "A plurality of RStores make up a logical disk that is presented to a user." EX1006, [0019]. "An RSEG is an atomic logical unit that is roughly analogous to a PSEG in the physical domain—akin to a logical disk partition of an RStore." EX1006, [0060]. In other words, a logical

disk is made up of logical units called RStores, which themselves are partitioned into logical units called RSEGs. *See* EX1006, [0060]-[0061], [0065]-[0070].

122. Lubbers describes that "[a]n RStore is analogous to a RAID disk set, differing in that it comprises PSEGs rather than physical disks." EX1006, [0068]. "An RStore is smaller than conventional RAID storage volumes, and so a given LUN 102 will comprise multiple RStores as opposed to a single RAID storage volume in conventional systems." EX1006, [0068]. Below is a diagram of Lubber's virtual storage structure utilized in Lubbers.

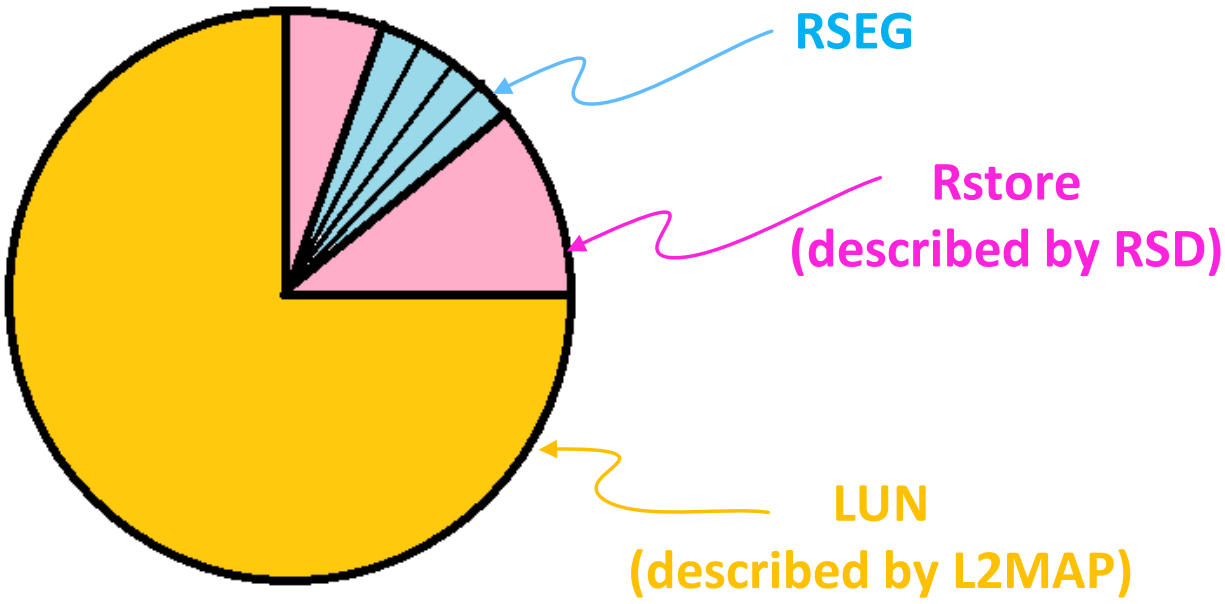


Diagram of Lubber's Virtual Storage Structure

123. As shown in Fig. 5 (reproduced below), Lubbers uses a set of data mapping tables to locate data within the logical and physical storage. For example,

"an **L2MAP 501**, an **LMAP 503**, and **a redundancy set descriptor (RSD) 505** as the primary structures for mapping a logical disk address to physical storage location(s) represented by that address." EX1006, [0057]. "A single L2MAP handles the entire LUN 102." EX1006, [0057]. Lubbers describes that "[t]he logical disk address specified in a request [] selects a particular entry within **LMAP 503** corresponding to a particular RSEG that in turn corresponds to 1 Mbyte address space allocated to the particular RSEG#." EX1006, [0060]. "Each LMAP entry also includes state information about the particular RSEG, and an RSD pointer." EX1006, [0060]. "The RSD pointer points to a specific **RSD 505** that contains metadata describing the RStore in which the corresponding RSEG exists." EX1006, [0062].

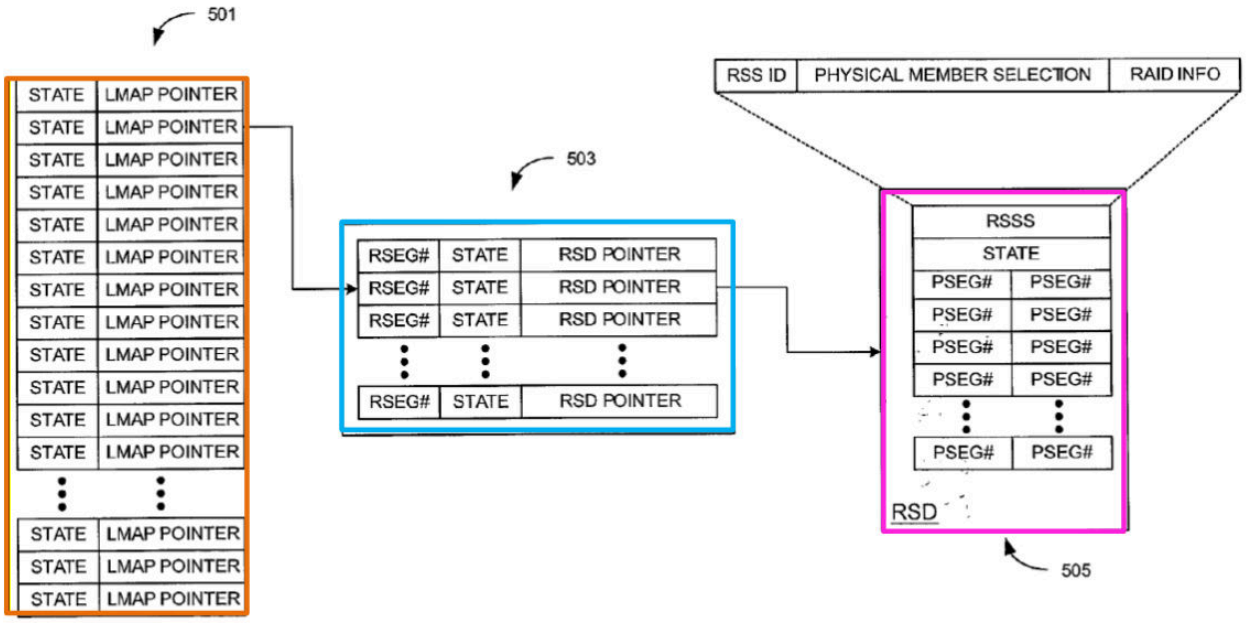


FIG. 5

EX1006, Fig. 5 (annotated)

124. "[E]ach request for storage access specifies a LUN 102, and an address." EX1006, [0063]. "The [logical disk address] LDA specified by the request is used to index into **L2MAP 501**, which in turn points to a specific one of the **LMAPs**." EX1006, [0063]. "The address specified in the request is used to determine an offset into the specified **LMAP** such that a specific RSEG that corresponds to the request-specified address is returned." EX1006, [0063]. "Once the RSEG# is known, the corresponding **RSD** is examined to identify specific PSEGs [physical storage segments] that are members of the redundancy segment, and metadata that enables a NSC 401 to generate drive specific commands to access the requested data." EX1006, [0063].

125. Lubber describes that "[t]he relationship of particular [physical] disks can be modified to provide various levels of protection and allow for construction of virtual disks with variable levels of data protection." EX1006, [0128]. "[M]ultiple types of RAID data protection may be implemented within a single LUN 102 such that a first range of logical disk addresses (LDAs) correspond to unprotected data, and a second set of LDAs within the same LUN 102 implement RAID 5 protection." EX1006, [0055]. "RAID information, or more generically data protection information, describes the type of data protection, if any, that is implemented in [a] particular RStore." "Hence, the data structures implementing

the memory representation must be flexible to handle this variety, yet efficient such that LUNs 102 do not require excessive data structures." EX1006, [0055]. "The RAID information, or more generically data protection information, describes the type of data protection, if any, that is implemented in the particular RStore." EX1006, [0062].

126. The NSC stores rules used to select the level of data protection. *See* EX1006, [0065]. A POSITA would have understood and found obvious that these rules would have been based on a number of factors, including, for example the type of application and its priority level, in order to better "satisfy business objectives." *See* EX1006, [0055]; EX1015, [0027] ("Application policies 16 concern an application's requirements on the data that it generates or uses—and, indirectly, the storage on which that data resides. For instance, an application may have specific requirements for the speed and format of data access, or for recoverability of the data in the event of an outage."), [0028] ("application attributes 10*c* may include initial access delay, sustained access rate, I/O rate, I/O bandwidth, read/write ratio, sequential/direct access, aggregate backup/recovery criteria"), [0031] (application policies "are designed for the management of storage (and data residing on the storage) and are intended to satisfy business objectives"), [0050] (a "data file may be defined at [] different levels in order to allocate specific storage resources to the data file"), [0051] (application level attributes "are

concerned with properties and characteristics related to the requirements demanded by the client application for data the application generates or uses"), [0053] (policies, like application policies, "for the management of storage, and data residing on the storage, may be created to satisfy business objectives," which may dictate levels of "reliability, performance, availability, and installability of [] storage elements"), [0054] (storage level attributes "may define such aspects as the number of storage devices, the RAID level, the use of backup or 'snapshot' programs, garbage collection, cache management, etc."). In other words, it would have been obvious that the rules implemented by Lubbers to select a data protection level would have been based, at least in part, on attributes of the application, including, for example, the priority of the application, as indicated by its requirements for the speed and format of data access and/or for recoverability of the data in the event of an outage.

127. In Lubbers, "RStores preferably implement a selected [level] of data protection." EX1006, [0019]. The membership of an RStore "is selected in a round-robin fashion such that the first RStore will have a PSEG selected from disk 0 as a member 0, a PSEG selected from Disk 1 as member 1, and so forth," according to a given data protection scheme. EX1006, [0079]. "Any number of RStores can be defined in this manner, and PSEGs are selected for any given RStore to meet the requirement of a given data protection scheme." EX1006,

[0079]. Accordingly, a logical address is associated with a given RStore, which in turn is structured and associated with a set of PSEGs based on the data protection requirements or rules of a given application.

128. Accordingly, Hashimoto-Lubbers ***select[s] one or more of the virtual assets*** (*e.g.*, the storage hypervisor of the storage system selects the RStore within a virtual disk associated with the logical disk address specified in an I/O command) ***for storage of the set of data*** (*e.g.*, when the I/O command is a write command, the I/O command results in storing data in the storage system) ***based on an allocation of data on the virtual storage assets*** (*e.g.*, RStores are associated with PSEGs and selected for inclusion within a LUN based on application-specific rules related to data protection level, and as a result the rules dictate how the data must be allocated on the virtual storage assets).

[1.4]

129. Lubber's "mapping structure that is implemented in memory of a NSC 401 [] enables translation of a request expressed in terms of a logical block address (LBA) from host 301 into a read/write command addressed to a particular portion of a physical disk 405 having the desired information." EX1006, [0054]. As discussed for [1.3], *supra*, "each request for storage access specifies a LUN 102, and an address." EX1006, [0063]. A NSC 401 of the storage system uses the data mapping tables (*e.g.*, L2MAP 501, LMAP 503, and RSD 505) "to generate drive

specific commands to access the requested data." EX1006, [0063].

130. Under Hashimoto-Lubbers, this process of address translation would have been implemented within the address translation described in Hashimoto. Specifically, Hashimoto describes that "[w]hen the virtual storage system (0) 240 receives a disk I/O command from the virtual machine (0) 130, the virtual storage system (0) 240 obtains the virtual disk number that is the target of the access by the disk I/O command to specify the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0229]. This translation happens according to the method described in Lubbers, discussed above. "Then, the storage hypervisor 214 accesses the physical disk drive 236 and reads/writes data." EX1005, [0229]. A POSITA would understand from the broad disclosure that Lubbers is indicating that it would store to any of the physical resources, and a POSITA would nevertheless find that obvious from the above disclosures in Lubbers.

131. Accordingly, Hashimoto-Lubbers ***caus[es] storage of the set of data at the selected virtual storage assets on any of the physical resources*** (*e.g.*, the storage hypervisor utilizes Lubbers partitioning and address translation techniques to determine how to store data being written as part of an I/O command on the

addressed virtual asset and any corresponding physical storage devices).

2. *Claim 2*

132. As discussed above with respect to [1.1.b], "storage system 200 can be roughly divided into a physical layer, a hypervisor layer, and a virtual storage layer." EX1005, [0104]. "The hypervisor layer is implemented by the storage hypervisors 214 and 224…" EX1005, [0106].

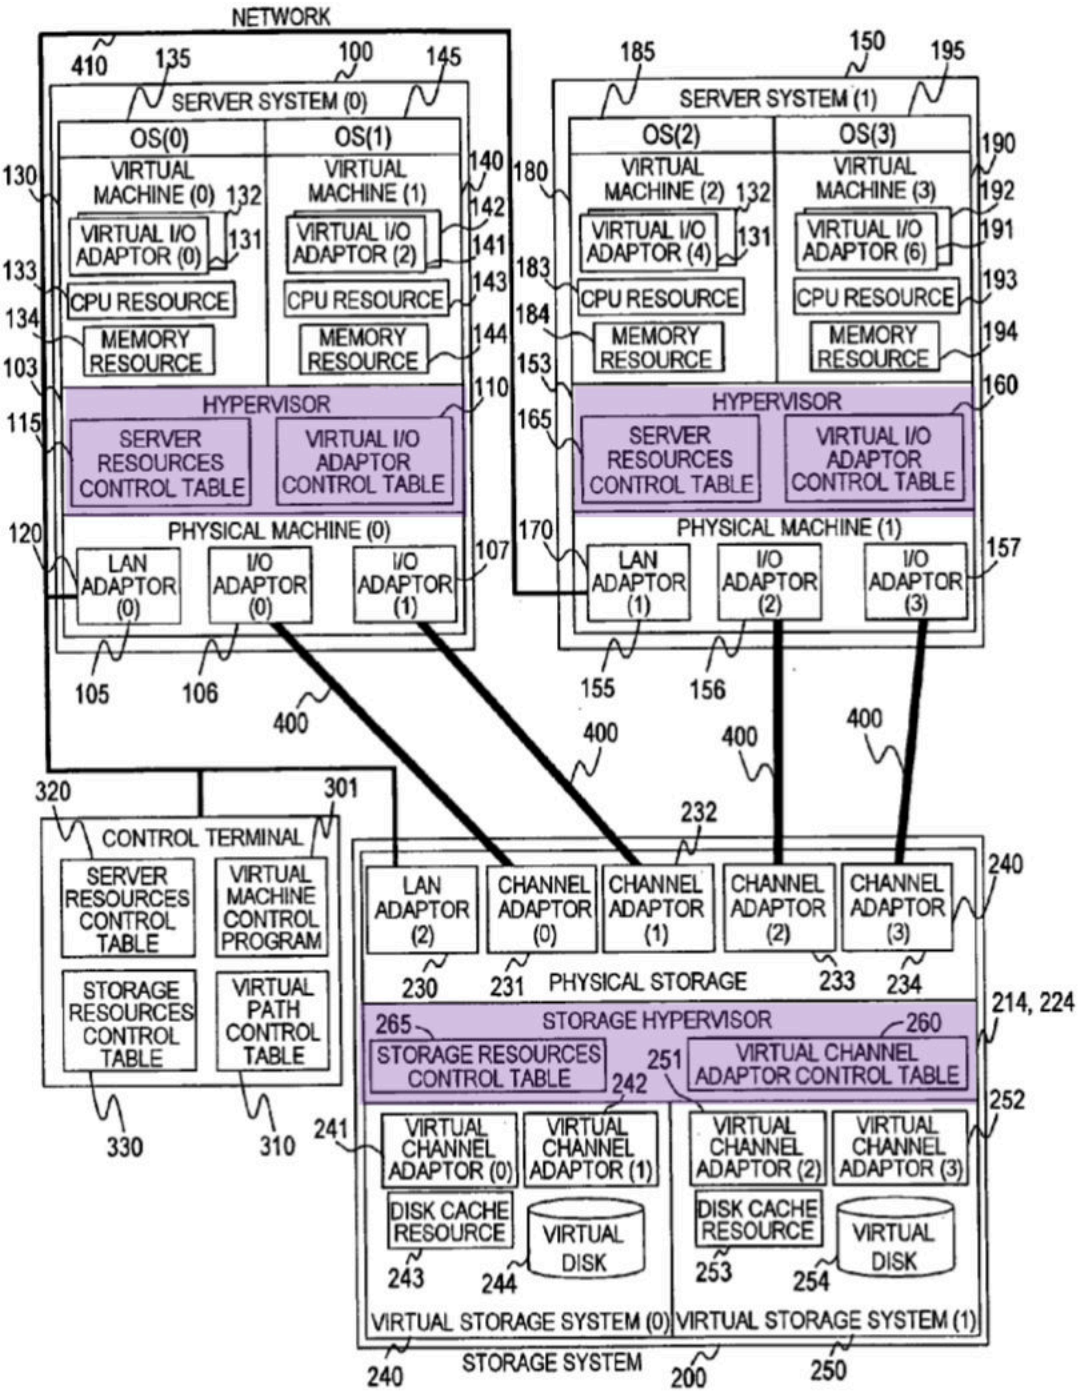


FIG.2

EX1005, FIG. 2 (Annotated)

133. The storage hypervisors 214 and 224 perform management functions including, for example, creating the virtual storage systems (*see* EX1005, [0184]-

[0185]), routing I/O requests generated by virtual machines (*see* EX1005, [0206], [0226]-[0227]), and detecting and fixing faults in I/O channels (*see* EX1005, [0247]-[0259]).

134. Accordingly, in Hashimoto-Lubbers, ***the at least one virtual storage asset*** (*e.g.*, virtual storage systems 240 and 250) ***is virtualized via storage virtualization layers running on each of the physical resources*** (*e.g.*, storage hypervisors 214 and 224 running on storage system 200).

3. *Claim 3*

135. As discussed above with respect to [1.1.a], "[t]he hypervisor 103, implemented by the execution of processing by the CPU (0) 101, is a means that enables logical partitioning of physical resources of the server system (0) 100." EX1006, [0059]. The "hypervisor 103 is a control program that creates logical partitions in the server system (0) 100 to create virtual machines that operate independently." EX1006, [0060]. These virtual machines are software instantiations of computer systems. *See* EX1005, [0004] ("[L]ogical partitioning technology can make one computer look like a plurality of virtual machines. When allocation of resources (processor, memory, etc.) to partitions is controlled, the performance of each virtual machine is assured."), [0062] ("The specification mostly describes operations chiefly in the software perspective, but, in reality, CPUs execute software to run the hypervisor 103 etc. and storage hypervisor 214

etc.").

136. Accordingly, in Hashimoto-Lubbers, ***at least two virtual assets*** (*e.g.*, virtual machines 130 and 140) ***each comprise software instantiations of computer systems***.

4. *Claim 4*

137. Hashimoto describes that the "virtual machine (0) 130 includes two virtual I/O adaptors (0) 131 and (1) 132." EX1005, [0164]. Similarly, the "virtual storage system (0) 240 has two virtual channel adaptors (0) 241 and (1) 242." EX1005, [0164]. "The virtual I/O adaptor (0) 131 is connected with the virtual channel adaptor (0) 241 and the virtual I/O adaptor (1) 132 is connected with the virtual channel adaptor (1) 242." EX1005, [0164]. "Virtual I/O adaptors are created by the hypervisor 103, 153, and virtual channel adaptors are created by the storage hypervisor 214, 224." EX1005, [0189].

138. The "virtual I/O adaptor (0) of the virtual machine (0) is implemented by server resources of the physical I/O adaptor (0)." EX1005, [0143]. In other words a virtual I/O adapter is an allocation of a physical I/O adaptor to a virtual machine. "The I/O adaptors 106 and 107 send data input/output requests and output data to the storage system 200 and receive data stored in the storage system 200." EX1005, [0057].

139. It was common for this functionality to at least in part be implemented

in a type of software called middleware. For example, "applications typically use a software component that operatively connects the application to [a] database allowing the application to store and retrieve the data it needs." EX1013, 1:20-23. "Since this software component is coupled between the application and the database, it is commonly referred to as middleware." EX1013, 1:23-25. In fact, Hashimoto describes that "application programs [running on the virtual machines] provide database services." EX1005, [0102]. This type of middleware was also present in the context of virtual machines. *See, e.g.*, EX1014, [0002], [0008], [0036]. Accordingly, it would have at least been obvious that, in order to implement the virtual I/O adapter, Hashimoto's virtual machines would run middleware to act as a bridge between a database (*e.g.*, stored in a virtual disk) and applications running on the virtual machine.

140. Accordingly, in Hashimoto-Lubbers, ***at least one of the virtual assets*** (*e.g.*, virtual machines 130, 140, 180, and 190) ***comprise middleware applications running on each of the physical resources*** (*e.g.*, at least a portion of the virtual I/O adapter functionality, which would be implemented as middleware) ***to access the virtual storage assets via the storage virtualization layer*** (*e.g.*, the virtual I/O adapter accesses a corresponding virtual storage system via the storage hypervisor and channel adapter of the virtual storage system).

5. *Claim 5*

141. Hashimoto describes that "[a] virtual machine control program 301 runs in the CPU of the control terminal 300." EX1005, [0083]. "The virtual machine control program 301 controls virtual configuration of the computer system by using the virtual path control table 310, the server resources control table 320, and the storage resources control table 330." EX1005, [0084]. For example, "the control terminal 300 includes a control unit (not shown) having connection information (the virtual path control table 310) that defines connection paths between virtual machines and virtual storage systems." EX1005, [0084].

142. When Hashimoto's system creates a virtual machine, "the administrator operates the control terminal 300 to set allocation of server resources of the server system (0) 100 and gives an instruction to create a virtual machine (S111)." EX1005, [0178]. Setting the allocation is a configuration of the virtual machine in relationship to the underlying physical resources. *See* EX1005, [0010], [0149]. "The control terminal 300 sends the resource allocation to the server system (0) 100." EX1005, [0178]. "The server system (0) 100 receives the virtual machine creation instruction from the control terminal 300 and then the hypervisor 103 creates a virtual machine by allocating server resources of the server system (0) 100, other than I/O adaptors, according to the received virtual machine creation instruction (S112)." EX1005, [0179].

143. Accordingly, in Hashimoto-Lubbers, ***the control center application is***

configured to perform configuration management or forecasting services via the virtualization layer instances (*e.g.*, the control terminal 300 and its virtual machine control program 301 perform configuration management via the hypervisors 103 and 153 and the storage hypervisors 214 and 224).

6. *Claim 7*

144. Hashimoto's "storage system 200 selects a single virtual disk or a plurality of virtual disks from the virtual disks 244 and 254 and provides the virtual disk or disks to the virtual machines 100 and 150 as storage areas." EX1005, [0115]. "The virtual disks thus selected are called logical units (LU)." EX1005, [0115]. These virtual disks are akin to the logical disks (also called logical units or LUNs) 102 described by Lubbers. *See* EX1006, [0034]

145. As discussed above with respect to [1.3], the LUNs 102 described by Lubbers are made up of a plurality of RStores, which is itself a logical unit. *See* EX1006, [0019], [0060] -[0061], [0065]-[0070]. Specifically, Lubbers describes that "[a]n RStore is analogous to a RAID disk set, differing in that it comprises PSEGs rather than physical disks." EX1006, [0068]. "An RStore is smaller than conventional RAID storage volumes, and so a given LUN 102 will comprise multiple RStores as opposed to a single RAID storage volume in conventional systems." EX1006, [0068]. As discussed above with respect to [1.3], Lubbers method for virtualizing a storage system selects RStores for inclusion in a LUN

based on the data protection level required for a given application, and as a result the data protection level dictates how the data must be allocated on the virtual storage assets. *See* EX1006, [0019], [0055], [0062], [0065], [0079].

146. Each LUN has "a unique identification," i.e., a logical unit number. EX1006, [0035]. ""Each LUN 102 is essential a contiguous range of logical addresses that can be addressed by host devices 105, 106, 107 and 109 by mapping requests from the connection protocol used by the hosts to the uniquely identified LUN 102." EX1006, [0035].

147. RStores are similarly uniquely identified within a LUN. *See* EX1006, [0037], [0042], [0057], [0062], [0065]-[0068]. Specifically, "[a] logical disk mapping layer maps a LDA specified in a request to a specific RStore as well as an offset within the RStore." EX1006, [0057]. An "RSD pointer points to a specific RSD 505 that contains metadata describing the RStore in which" corresponding RSEGs exist. EX1006, [0062]. In other words, each RSD uniquely identifies a corresponding RStore. *See* EX1006, [0057], [0062].

148. Accordingly, in Hashimoto-Lubbers, ***the at least one virtual storage asset*** (*e.g.*, the RStores—which are "logical units" and "analogous to a RAID disk set"—that are selected for inclusion in the one or more LUNs that are provided to a virtual machine) ***is a logical unit with a logical unit number (LUN)*** (*e.g.*, a "uniquely identified LUN" is associated with one or more RSDs that uniquely

identify the RStores that make up the LUN).

7. *Claim 8*

[8.1]

149. As discussed for [1.2], *supra*, FIG. 14 of Hashimoto "is a flowchart of a disk I/O command transmission process." EX1005, [0222]. As part of this process, the OS of a virtual machine issues an I/O command, which is ultimately communicated to a corresponding virtual storage system via a hypervisor and a storage hypervisor. *See* EX1006, [0222]-[0227]. The I/O command "specif[ies] the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229].

150. "When the virtual storage system (0) 240 receives a disk I/O command from the virtual machine (0) 130, the virtual storage system (0) 240 obtains the virtual disk number that is the target of the access by the disk I/O command to specify the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0229]. That is, the I/O commands are able to access (*e.g.*, read) any data within the virtual storage system to which the issuing virtual machine has access, which would include data previously written by the virtual machine to the virtual storage system.

See EX1005, [0183] ("when a virtual machine is created for the first time in the server system (0) 100, it is necessary to create a virtual storage system that the virtual machine accesses"); *see also id.*, [0116] ("[t]he virtual path control table 310 describes relations between the virtual machines and the virtual storage systems"), [0168] ("FIG. 8 shows the relations among the virtual machine numbers 316, the virtual I/O adaptor numbers 315, the virtual channel adaptor numbers 312, and the virtual storage system numbers 311 in the virtual path control table 310"), [0196] ("[t]he control terminal 300 is now capable of managing the relation from the virtual storage system to the virtual machine"), [0205] ("[t]he virtual machine (0) 130 issues that I/O command to the virtual storage system (0) 240" and "[t]he virtual machine (0) 130 issues that I/O command to the virtual storage system (0) 240")

151. Accordingly, Hashimoto-Lubbers ***obtain[s] a request from a user device to obtain at least a portion of the set of stored data*** (*e.g.*, an I/O request that reads data from a corresponding virtual disk).

[8.2]

152. As discussed above with respect to [8.1], the OS of a virtual machine issues an I/O command, which is ultimately communicated to a corresponding virtual storage system via a hypervisor and a storage hypervisor. *See* EX1006, [0222]-[0227]. The I/O command "specif[ies] the accessed virtual disk and

reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0229] "Then, the storage hypervisor 214 accesses the physical disk drive 236 and reads/writes data." EX1005, [0229]..

153. Accordingly, Hashimoto-Lubbers ***retriev[es] at least the portion of the set of stored data from the virtual storage asset*** (*e.g.*, the storage hypervisor 214 accesses the physical disk drive 236 and reads data).

[8.3]

154. Hashimoto describes that "[w]hen completing the data read from…the physical disk, the virtual storage system (0) 240 issues a disk I/O completion notification." EX1005, [0231]. "The storage hypervisor 214 then adds information to the disk I/O completion notification (see FIG. 12) and transfers the disk I/O completion notification to the physical channel adaptor (0) 231 (S142)." EX1005, [0232]. FIG. 12 shows that the disk I/O layer information includes data. *See* EX1005, FIG. 12, [0212]-[0213]. Indeed, in the case of a read I/O command, the disk I/O completion notification would naturally include the data being read. *See* EX1005, [0212] ("[t]he hypervisor communication payload 1204 is data transmitted in the hypervisor communication layers 1101 and 1105"), [0213] (with respect to Fig. 12, the hypervisor communication payload in communications

between a virtual machine and virtual storage system "includes a disk I/O command and transferred data"), [0229] ("the virtual storage system (0) 240…reads[] data []from the corresponding virtual disk 244").

155. "The channel adaptor (0) 231 [of storage system 200] sends the disk I/O completion notification to the connected I/O adaptor (0) 106 (the virtual machine (0) 130) (S143)." EX1005, [0233]. "The I/O channels 400 are formed of a transmission medium that can make communication according to a protocol suitable for data transfer, e.g., the Fibre Channel protocol." EX1005, [0086]. "In the first embodiment, the server systems 100 and 150 and the storage system 200 are connected in a one-to-one manner, but they may be connected through a network (SAN)…" EX1005, [0086].

156. Accordingly, Hashimoto-Lubbers ***transmit[s] at least the portion of the set of stored data to the user device*** (*e.g.*, the read data is contained in a disk I/O completion notification transmitted from storage system 200 to a corresponding virtual machine on server system 100, 150) ***via a network*** (*e.g.*, via a SAN interconnecting server systems 100 and 150 and the storage system 200).

B. Ground 2: Claims 9-20 Are Obvioius Over The Combination of Hashimoto, Lubbers, and Hirschfeld

1. *Claim 9*

[9.pre]

157. In my opinion, the Hashimoto–Lubbers–Hirschfeld combination

provides [9.pre] for substantially the same reasons discussed for [1.pre], which are unaffected by the addition of Hirschfeld.

[9.1]

158. As discussed for [1.1.a], *supra*, Hashimoto describes a computer system that includes server systems 100 and 150, at least one storage system 200, and a control terminal 300, all of which are connected via a network 410. EX1005, [0052]; *see also id.*, [0371]-[0375]. "A virtual machine control program 301 runs in the CPU of the control terminal 300." EX1005, [0083]. "The virtual machine control program 301 controls virtual configuration of the computer system by using the virtual path control table 310, the server resources control table 320, and the storage resources control table 330." EX1005, [0084]. "The control terminal 300 manages information for controlling the computer system and controls operations of the entire computer system." EX1005, [0052]. This virtual machine control program 301 running on control terminal 300 performs many of the same control function as Hirschfeld's SCM. *See* EX1007, [0030] ("virtualized logical server cloud 101 includes the SCM 102 that interfaces and controls each of the logical servers 106 and the bank of physical servers 108"). Thus, in Hashimoto-Lubbers-Hirschfeld, an SCM/control terminal would perform the described control functions.

159. Accordingly, Hashimoto-Lubbers-Hirschfeld includes ***a control***

computer running a control center application (*e.g.*, SCM/control terminal running a virtual machine control program).

[9.2.a]

160. The combination provides [9.2.a] for substantially the same reasons discussed for [1.1.a], which are unaffected by the addition of Hirschfeld.

[9.2.b]

161. The combination provides [9.2.b] for substantially the same reasons discussed for [1.1.b], which are unaffected by the addition of Hirschfeld.

[9.2.c]

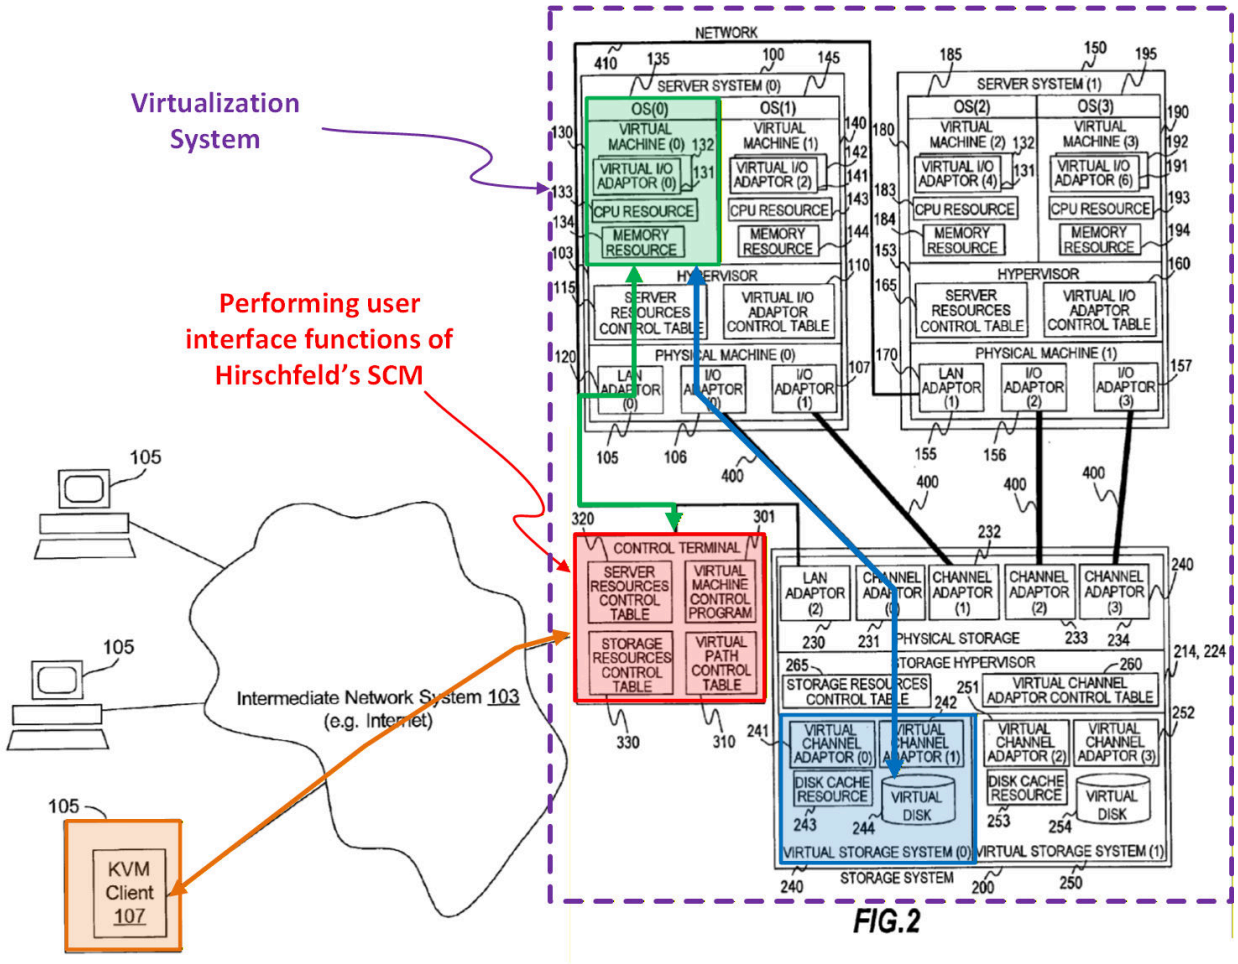
162. The combination provides [9.2.c] for substantially the same reasons discussed for [1.1.c], which are unaffected by the addition of Hirschfeld.

[9.3]

163. As discussed for [8.1], FIG. 14 of Hashimoto "is a flowchart of a disk I/O command transmission process." EX1005, [0222]. As part of this process, the OS of a virtual machine issues an I/O command, which is ultimately communicated to a corresponding virtual storage system via a hypervisor and a storage hypervisor. *See* EX1006, [0222]-[0227]. The I/O command "specif[ies] the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229].

164. As discussed for [9.1], *supra*, in Hashimoto-Lubbers-Hirschfeld, the virtual machine is controlled via the SCM/control terminal. *See* EX1005, [0030],

[0052], [0083]-[0084], [0222]-[0227]; EX1007, [0030]-[0032]. For example, Hirschfeld describes that "the SCM 102 that interfaces and controls each of the logical servers 106 and the bank of physical servers 108." EX1007, [0030]. Additionally, "the SCM 102 includes a user interface (I/F) 104 that provides a web page portal or other access mechanism to facilitate communication with the remote terminals 105 via the network system 103." EX1007, [0030]. "In this manner, a user of any remote terminal 105 may login to and access any of the logical servers 106 of the virtualized logical server cloud 101 via the user interface 104 and SCM 102 as though logging into a server directly or via a local area network (LAN)." EX1007, [0030]. "The graphical user interface (GUI) of the user interface 104 provides the mechanism (portal or other interface) through which the user may access one or more logical servers 106 based on permission," and "control a selected logical server." EX1007, [0031].



Amalgamation of EX1005 Fig. 2 and EX1007 Fig. 1

165. Thus, as shown in the above diagram of Hashimoto-Lubbers-Hirschfeld, a user would control one or more of the virtual machines 130, 140, 180, and 190 via a remote terminal in communication with the combination's SCM/control terminal, which in turn passes the user's commands onto the server system(s) hosting the virtual machine(s) to which the user has access. *See* EX1007, [0029]-[0032]. For example, Hashimoto describes that the virtual machines run "application programs [that] provide…services…to clients (not

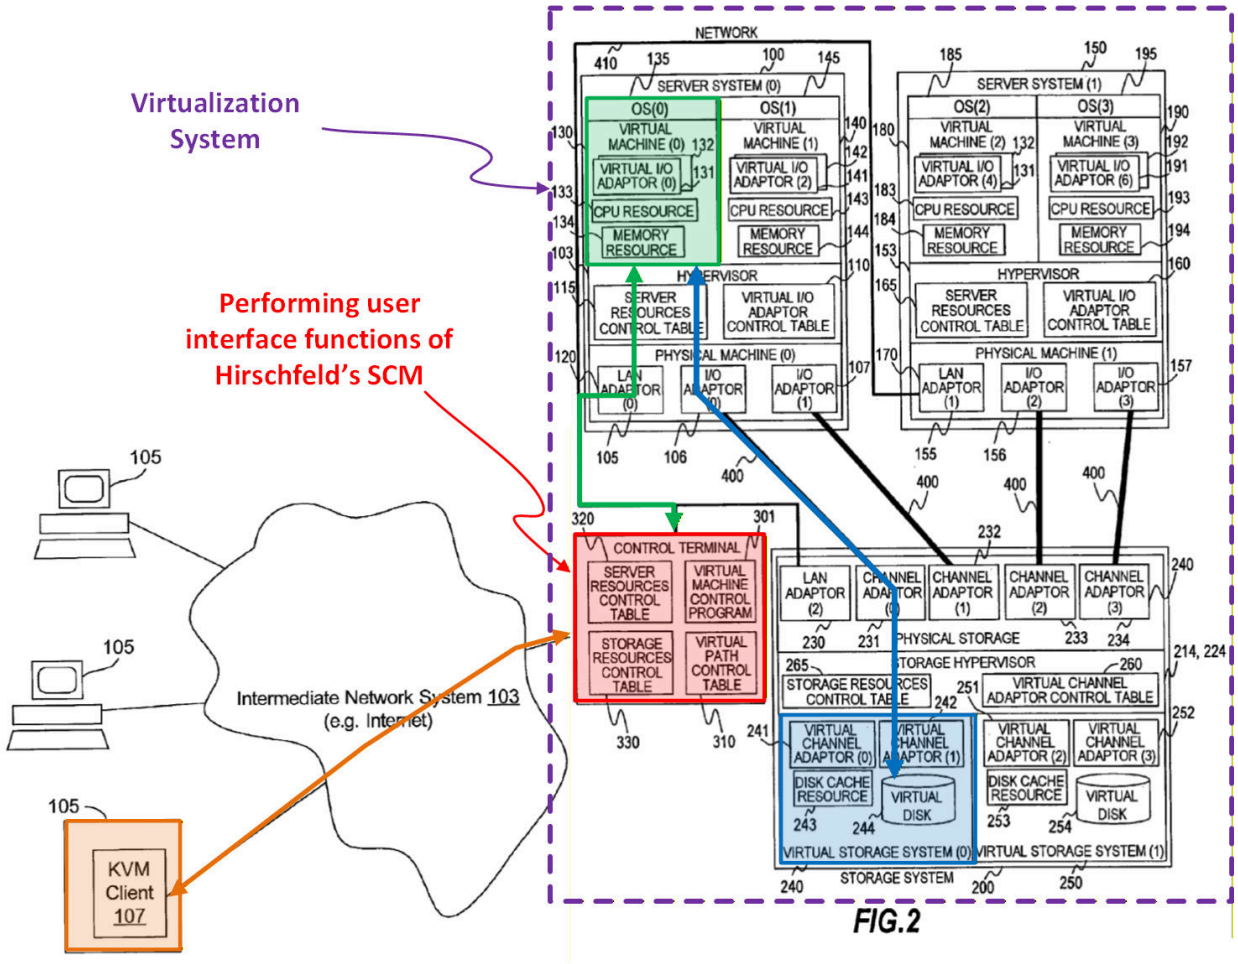
shown) connected to the server system (0) 100." EX1005, [0102]. In Hashimoto-Lubbers-Hirschfeld, the user would interact with and control these application programs via the SCM/control terminal. *See* EX1007, [0029]-[0032]. Hasimoto describes that the "application programs provide database services, Web services, etc." EX1005, [0102]. "[F]rom web servers, database servers, spreadsheets and word-processing programs, to games and browsers," there are many well-known application programs run on virtual machines. EX1012, 11:50-53. In the case of a spreadsheet or word-processing program, it was well-known that a user could open a spreadsheet or document previously stored in the storage system. Similarly, in the case of a game, it was well-known that starting the game requires loading the components of the game stored in an associated storage system.

166. Accordingly, in Hashimoto-Lubbers-Hirschfeld, ***the control center application*** (*e.g.*, the control application run by the SCM/control terminal of the combination implementing the user interface functionality of Hirschfeld's SCM) ***is configured to obtain a request from a user device for data stored at any of the virtual storage assets*** (*e.g.*, a user provides commands—including, for example, "open" commands that retrieve previously stored documents from a virtual storage system—through their remote terminal that are transmitted to the SCM/control terminal, which then pass these commands to the virtual machine with which the user is interfacing).

[9.4]

167. As discussed for [9.3], Hirschfeld describes that "a user of any remote terminal 105 may login to and access any of the logical servers 106 of the virtualized logical server cloud 101 via the user interface 104 and SCM 102 as though logging into a server directly or via a local area network (LAN)." EX1007, [0030]. "[A] terminal 105 executes a keyboard/video/mouse (KVM) client 107 that transmits keyboard and mouse commands to a selected logical server 106 and that receives video and audio information from the selected logical server 106." EX1007, [0031]. In other words, the output associated with commands issued by the user are provided via the SCM 102.

168. As discussed for [8.2], Hashimoto describes that the OS of a virtual machine issues an I/O command, which is ultimately communicated to a corresponding virtual storage system via a hypervisor and a storage hypervisor. *See* EX1006, [0222]-[0227]. The I/O command "specif[ies] the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0229] "Then, the storage hypervisor 214 accesses the physical disk drive 236 and reads/writes data." EX1005, [0229].



Amalgamation of EX1005 Fig. 2 and EX1007 Fig. 1

169. Thus, as shown in the above diagram of Hashimoto-Lubbers-Hirschfeld, a user issues a command from their remote terminal 105, which **transmits the command to the combination's SCM/control terminal**. *See* EX1007, [0030]-[0031]. As discussed for [9.3], *supra*, this command may be with respect to an application program being run by the virtual machine with which the user is interacting. *See* EX1005, [0105]; EX1007, [0029]-[0032]. The **SCM/control terminal passes the user's command to a virtual machine 130,**

140, 180, or 190 via network 410. *See* EX1007, [0030]-[0031]; EX1005, [0052], [0083]-[0087]. When the user's command requires the application program to access data stored in memory (*e.g.*, a request to "open" a spreadsheet or document), the application program interfaces with the virtual machines OS to access the stored data. *See* EX1005, [0101]-[0102].

170. To access stored data, **the OS of the virtual machine 130 transmits the I/O command to a corresponding virtual storage system 240 or 250**, consistent with the teachings of Hashimoto. *See* EX1005, Fig. 14, [0222]-[0227]. As discussed for [8.3], *supra*, in the case of a read I/O command, a disk I/O completion notification that would naturally include the data being read (*e.g.*, the spreadsheet or document being opened by an application program) would be returned to the virtual machine 130 that issued the I/O command. *See* EX1005, Fig. 15, [0212]-[0213], [0228]-[0235].

171. Accordingly, in Hashimoto-Lubbers-Hirschfeld, ***the control center application*** (*e.g.*, the control application run by the SCM/control terminal of the combination) ***is configured to retrieve the data from one or more virtual storage assets that are storing the data*** (*e.g.*, the SCM/control terminal of the combination passes a user's application program commands to the virtual machine with which the user is interfacing, which in turn transmits corresponding I/O commands to a corresponding virtual storage system, and display of the result of the I/O

commands is sent back to the SCM/control terminal by the virtual machine on behalf of the application program).

[9.5]

172. As discussed for [9.4], Hirschfeld describes that "a user of any remote terminal 105 may login to and access any of the logical servers 106 of the virtualized logical server cloud 101 via the user interface 104 and SCM 102 as though logging into a server directly or via a local area network (LAN)." EX1007, [0030]. "[A] terminal 105 executes a keyboard/video/mouse (KVM) client 107 that transmits keyboard and mouse commands to a selected logical server 106 and that receives video and audio information from the selected logical server 106." EX1007, [0031]. In other words, the output associated with the commands issued by the user to the application program running on a virtual machine is sent back to the user's terminal 105 via SCM 102.

173. Accordingly, in Hashimoto-Lubbers-Hirschfeld, ***the control center application is configured to transmit the retrieved data to the user device***.

2. *Claim 10*

174. Hashimoto–Lubbers–Hirschfeld provides [10] for substantially the same reasons discussed for [2], which are unaffected by the addition of Hirschfeld.

3. *Claim 11*

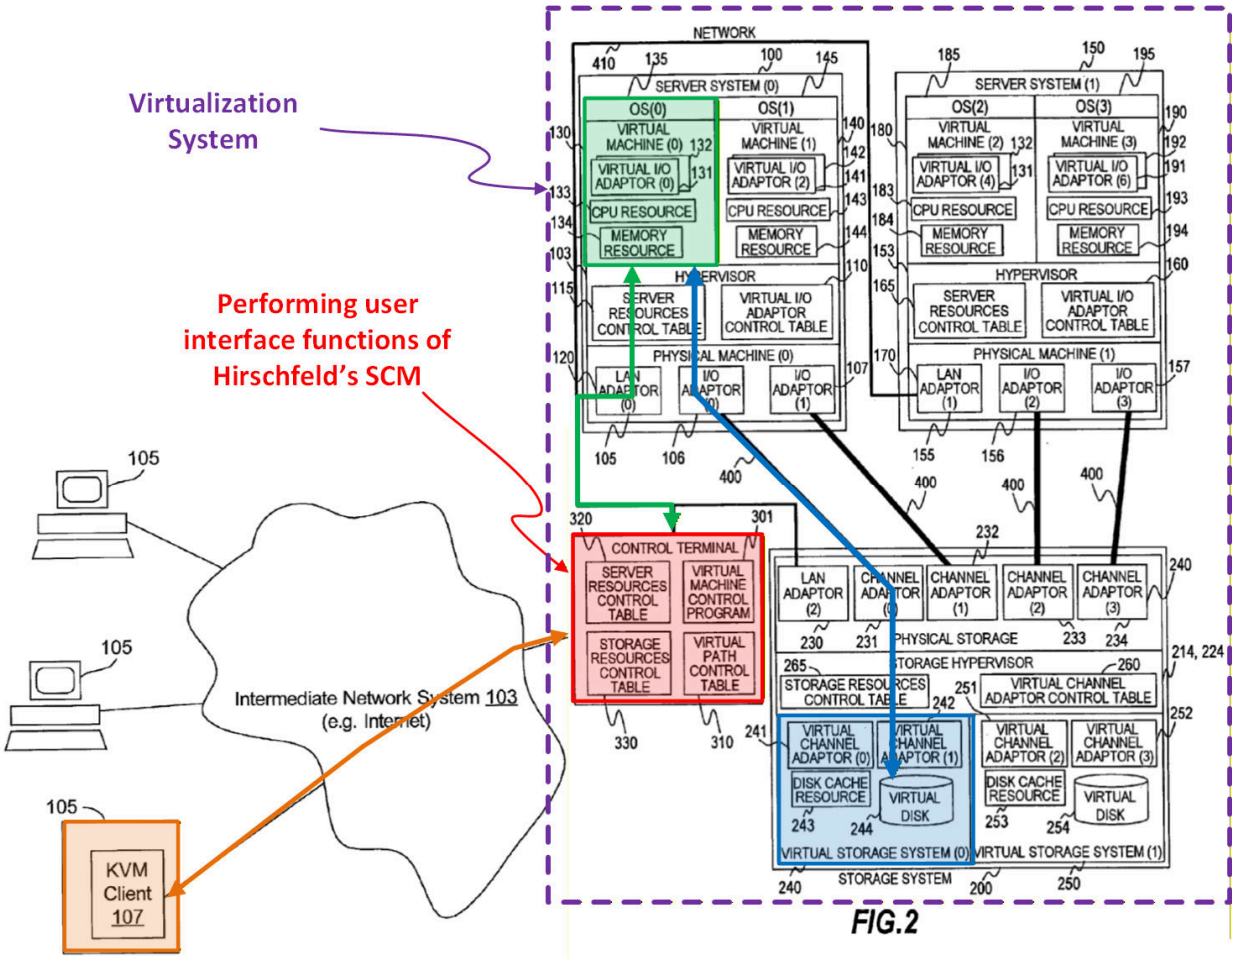
[11.1]

175. Initially, the phrase "the data to be stored" lacks antecedent basis in claim 9. *See* EX1001, 24:55–25:3. Claim 9 recites "a request … for data stored" and "retriev[ing] the data from one or more virtual storage assets" (akin to a read I/O). EX1001, 24:66–25:2; *see* EX1005, [0229]-[0231]. By contrast, "obtain[ing] … data to be stored" and "stor[ing] the data across any of the virtual storage assets" is akin to a write I/O. *See* EX1001, 25:7–12; EX1005, [0229]-[0231]. For purposes of this IPR only, where §112 grounds are unavailable, I assume [11.1] concerns a separate request to "store … data," distinct from the "request … for data stored" recited in [9.3].

176. As discussed for [9.4], Hirschfeld describes that "a user of any remote terminal 105 may login to and access any of the logical servers 106 of the virtualized logical server cloud 101 via the user interface 104 and SCM 102 as though logging into a server directly or via a local area network (LAN)." EX1007, [0030]. "[A] terminal 105 executes a keyboard/video/mouse (KVM) client 107 that transmits keyboard and mouse commands to a selected logical server 106 and that receives video and audio information from the selected logical server 106." EX1007, [0031]. In other words, the output associated with the commands issued by the user to the application program running on a virtual machine is sent back to the user's terminal 105 via SCM 102.

177. As discussed for [8.2], Hashimoto describes that the OS of a virtual

machine issues an I/O command, which is ultimately communicated to a corresponding virtual storage system via a hypervisor and a storage hypervisor. *See* EX1006, [0222]-[0227]. The I/O command "specif[ies] the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0229] "Then, the storage hypervisor 214 accesses the physical disk drive 236 and reads/writes data." EX1005, [0229].



Amalgamation of EX1005 Fig. 2 and EX1007 Fig. 1

178. Thus, as shown in the above diagram of Hashimoto-Lubbers-Hirschfeld, a user would control one or more of the virtual machines 130, 140, 180, and 190 via **a remote terminal in communication with the combination's SCM/control terminal**, which in turn **passes the user's commands onto the server system(s) hosting the virtual machine(s) to which the user has access**. *See* EX1007, [0029]-[0032]. As discussed for [9.3], *supra*, Hashimoto describes that the virtual machines run "application programs [that] provide…services…to clients (not shown) connected to the server system (0) 100." EX1005, [0102]. In Hashimoto-Lubbers-Hirschfeld, the user would interact with and control these application programs via the SCM/control terminal. *See* EX1007, [0029]-[0032]. In the case of a spreadsheet or word-processing program, it was well-known that a user could enter data into a document or spreadsheet file and "save" the file, which would then be stored on an associated storage system.

179. **The SCM/control terminal passes the user's command to a virtual machine 130 via network 410**. *See* EX1007, [0030]-[0031]; EX1005, [0052], [0083]-[0087]. When the user's command requires the application program to store data in memory (*e.g.*, a request to "save" a spreadsheet or document file created through user input), the application program interfaces with the virtual machines OS to store the data. *See* EX1005, [0101]-[0102]. Specifically, **the OS**

of the virtual machine 130 transmits a write I/O command to a corresponding virtual storage system 240 or 250, consistent with the teachings of Hashimoto. *See* EX1005, Fig. 14, [0222]-[0227]

180. Accordingly, Hashimoto-Lubbers-Hirschfeld ***obtain[s] the data to be stored at any of the virtual storage assets*** (*e.g.*, the SCM/control terminal receives data input by a user with the KVM client 107 of remote terminal 105).

[11.2]

181. As discussed for [9.4], Hashimoto describes that "[w]hen the virtual storage system (0) 240 receives a disk I/O command from the virtual machine (0) 130, the virtual storage system (0) 240 obtains the virtual disk number that is the target of the access by the disk I/O command to specify the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0229].

182. In order to translate the address, Lubber teaches that a "mapping structure…enables translation of a request expressed in terms of a logical block address (LBA) from host 301 into a read/write command addressed to a particular portion of a physical disk 405 having the desired information." EX1006, [0054]. As discussed for [1.3], *supra*, "each request for storage access specifies a LUN

102, and an address." EX1006, [0063]. In Hashimoto-Lubbers-Hirschfeld, the hypervisor 214 of the corresponding virtual storage system uses storage mapping tables (e.g., L2MAP 501, LMAP 503, and RSD 505) "to generate drive specific commands to access the requested data." EX1006, [0063].

183. Thus, in the example discussed with respect to [11.1], the word processing application and OS being run by the virtual machine would have issued an I/O command to write the document input by the user into the virtual storage system.

184. Accordingly, Hashimoto-Lubbers-Hirschfeld ***store[s] the data across any of the virtual storage assets running on the physical resources*** (*e.g.*, the SCM/control terminal receives an input associated with a save command from a user via the KVM client 107 of remote terminal 105, and forwards this input to the virtual machine, which in turn issues a write I/O command to the virtual storage system).

4. *Claim 12*

185. Hashimoto–Lubbers–Hirschfeld provides [12] for substantially the same reasons discussed for [3], which are unaffected by the addition of Hirschfeld.

5. *Claim 13*

186. Hashimoto–Lubbers–Hirschfeld provides [13] for substantially the

same reasons discussed for [4], which are unaffected by the addition of Hirschfeld.

6. *Claim 14*

187. Hashimoto–Lubbers–Hirschfeld provides [14] for substantially the same reasons discussed for [5], which are unaffected by the addition of Hirschfeld.

7. *Claim 15*

188. Hashmimoto describes that "[a]pplication programs run on the OSs 135 and 145" of virtual machines 130 and 140. EX1005, [0101]-[0102]. "These application programs provide database services, Web services, etc. to clients (not shown) connected to the server system (0) 100." EX1005, [0102]. Hashimoto teaches that "when a virtual machine is created for the first time in the server system (0) 100, it is necessary to create a virtual storage system that the virtual machine accesses." EX1005, [0183]. In other words, the virtual storage system is the resource a virtual machine uses for storage, and thus the virtual storage system is an obvious place that the virtual machine would store one or more databases associated with the "database services" application program run by the virtual machine. *See* EX1005, [0101]-[0102], [0183].

189. Accordingly, in Hashimoto-Lubbers-Hirschfeld, ***the virtual storage assets comprise databases***.

8. *Claim 16*

190. Hashimoto–Lubbers–Hirschfeld provides [16] for substantially the same reasons discussed for [7], which are unaffected by the addition of Hirschfeld.

9. *Claim 17*

191. Hashimoto describes that "[t]he administrator can manage the computer system by operating the control terminal 300." EX1005, [0310], [0321]; *see also id.*, [0052] ("The control terminal 300 manages information for controlling the computer system and controls operations of the entire computer system"). For example, "[t]he administrator specifies allocation of the storage resources by means of the control terminal 300 and the storage resources control table 265 is generated as per the administrator's setting." EX1005, [0160]. These allocation instructions are communicated to the server and storage systems via network 410. EX1005, [0184] ("the control terminal 300 sends allocation of storage resources of the storage system 200 to the storage system 200 and gives an instruction to create a virtual storage system"); *see also id.*, [0085] ("[t]he LAN interface [of the control terminal 300] sends and receives control signals and control information (the contents of the various control tables) to and from the systems connected through the network 410.").

192. Accordingly, in Hashimoto-Lubbers-Hirschfeld, ***the virtualization***

system is further configured to receive, over a network by a user, one or more instructions to configure the virtual storage assets (*e.g.*, the storage system 200 receives control instructions to configure virtual storage systems 240 and 250 over network 410 from control terminal 300 as a result of input from an administrator—a type of user).

10. *Claim 18*

[18.pre]

193. Hashimoto–Lubbers–Hirschfeld provides [18.pre] for substantially the same reasons discussed for [1.pre], which are unaffected by the addition of Hirschfeld.

[18.1]

194. The combination includes ***a virtualization system that includes the control center application, the one or more physical resources, and one or more virtual assets on each of the physical resources*** for substantially the same reasons discussed for [1.1.a], *supra*, which are unaffected by the addition of Hirschfeld.

195. As discussed for [9.3], in Hashimoto-Lubbers-Hirschfeld, control of the virtual machine is controlled via the SCM/virtual machine. *See* EX1005, [0030], [0052], [0083]-[0084], [0222]-[0227]; EX1007, [0030]-[0032]. For example, Hirschfeld describes that "the SCM 102 that interfaces and controls each of the logical servers 106 and the bank of physical servers 108." EX1007, [0030].

Additionally, "the SCM 102 includes a user interface (I/F) 104 that provides a web page portal or other access mechanism to facilitate communication with the remote terminals 105 via the network system 103." EX1007, [0030]. "In this manner, a user of any remote terminal 105 may login to and access any of the logical servers 106 of the virtualized logical server cloud 101 via the user interface 104 and SCM 102 as though logging into a server directly or via a local area network (LAN)." EX1007, [0030]. "The graphical user interface (GUI) of the user interface 104 provides the mechanism (portal or other interface) through which the user may access one or more logical servers 106 based on permission," and "control a selected logical server." EX1007, [0031]. Thus, in Hashimoto-Lubbers-Hirschfeld, a user would control one or more of the virtual machines 130, 140, 180, and 190 via a remote terminal in communication with the combination's SCM/control terminal, which in turn passes the user's commands onto the server system(s) hosting the virtual machine(s) to which the user has access.

196. As discussed for [8.2], Hashimoto describes that the OS of a virtual machine issues an I/O command, which is ultimately communicated to a corresponding virtual storage system via a hypervisor and a storage hypervisor. *See* EX1006, [0222]-[0227]. The I/O command "specif[ies] the accessed virtual disk and reads/writes data to/from the corresponding virtual disk 244." EX1005, [0229]. "In this process, the storage hypervisor 214 translates the virtual block

address of the accessed virtual disk 244 to a physical block address of the physical disk drive 236." EX1005, [0229] "Then, the storage hypervisor 214 accesses the physical disk drive 236 and reads/writes data." EX1005, [0229].

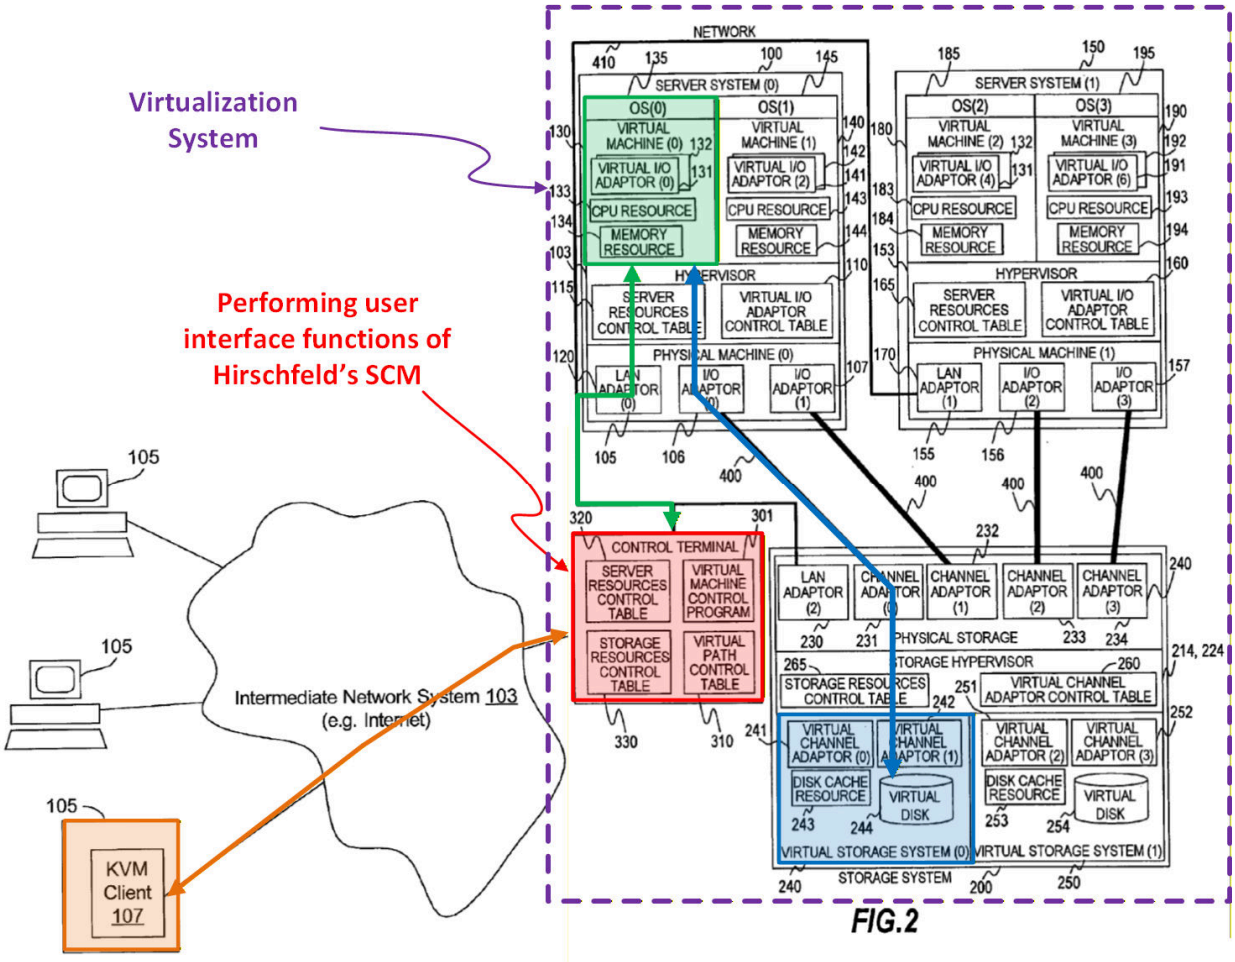


FIG.2

Amalgamation of EX1005 Fig. 2 and EX1007 Fig. 1

197. Thus, as shown in the above diagram of Hashimoto-Lubbers-Hirschfeld, a user issues a command from their remote terminal 105, which **transmits the command to the combination's SCM/control terminal**. *See* EX1007, [0030]-[0031]. As discussed for [9.3], *supra*, this command may be with

respect to an application program being run by the virtual machine with which the user is interacting. *See* EX1005, [0105]; EX1007, [0029]-[0032]. In the case of a spreadsheet or word-processing program, it was well-known that a user could enter data into a document or spreadsheet file and "save" the file, which would then be stored on an associated storage system.

198. **The SCM/control terminal passes the user's command to a virtual machine 130, 140, 180, or 190** via network 410. *See* EX1007, [0030]-[0031]; EX1005, [0052], [0083]-[0087]. When the user's command requires the application program to store data in memory (*e.g.*, a request to "save" a spreadsheet or document file created through user input), the application program interfaces with the virtual machines OS to store the data. *See* EX1005, [0101]-[0102]. Specifically, **the OS of the virtual machine 130 transmits a write I/O command to a corresponding virtual storage system 240 or 250**, consistent with the teachings of Hashimoto. *See* EX1005, Fig. 14, [0222]-[0227].

199. Accordingly, Hashimoto-Lubbers-Hirschfeld ***receiv[es] a request to store data, at a control center application*** (*e.g.*, the SCM/control terminal receives an I/O command from the KVM client of a remote terminal) ***in communication with one or more physical resources that are part of [the] virtualization system*** (*e.g.*, the SCM/control terminal is connected to one or more server systems and one or more storage systems).

[18.2]

200. Hashimoto–Lubbers–Hirschfeld provides [18.2] for substantially the same reasons discussed for [1.3], which are unaffected by the addition of Hirschfeld.

[18.3]

201. Hashimoto–Lubbers–Hirschfeld provides [18.3] for substantially the same reasons discussed for [11.2].

[18.4]

202. Hashimoto–Lubbers–Hirschfeld provides [18.4] for substantially the same reasons discussed for [2], which are unaffected by the addition of Hirschfeld.

11. *Claim 19*

203. Hashimoto–Lubbers–Hirschfeld provides [19] for substantially the same reasons discussed for [1.1.a] and [1.1.b], which are unaffected by the addition of Hirschfeld.

12. *Claim 20*

204. Hashimoto–Lubbers–Hirschfeld provides [20] for substantially the same reasons discussed for [3] and [15].

C. Ground 3: Claims 6 Are Obvioius Over The Combination of Hashimoto, Lubbers, and Khandekar

1. *Claim 6*

205. As part of deploying VMs, Khandekar teaches that "the user is given some choice about how the provisioned VM's virtual hardware, system software, and applications are selected." EX1012, 11:11-13. "If the user's configuration preferences cannot be matched exactly, then [a] heuristics engine 880 [of the provisioning server] suggests the best available match by applying a database of heuristic rules about properties of VMs, available applications and operating systems, virtualized devices, etc." EX1012, 11:14-19. The heuristics engine 880 also selects an appropriate host computer on which to deploy the VM and its applications "in order to balance the load on the different available hosts." EX1012, 13:4-8.

206. Khandekar describes that, "[o]nce a VM has been configured and deployed, note that it will also be possible to alter its configuration, for example by upgrading the software and/or hardware, installing different applications, etc." EX1012, 4:48-51. "[M]aintenance tasks like upgrade of software and hardware (for example, activating a different number of processors, memory chips, etc.) installation of service packs or hot fixes, detection of virus, etc., can be run by the administrative component of the provisioning server." EX1012, 13:46-51. "Hosts can be monitored for load and if the load crosses a threshold, the administrator can be alerted and/or actions can be taken such as migrating VMs to a lightly loaded server." EX1012, 13:51-54. In other words, Khandekar describes dynamic

deployment or allocation of applications to virtual assets based on business policies (*e.g.*, upgrades and hot fixes) and workload conditions (*e.g.*, the load of the hosts/servers).

207. In Hashimoto-Lubbers-Khandekar, the above-describe deployment and management functionality would be carried out by the control terminal 300, which "controls operations of the entire computer system," including "set[ting] allocation of server resources of the server system (0) 100 and giv[ing] an instruction to create a virtual machine." EX1005, [0052], [0084], [0178].

208. Accordingly, in Hashimoto-Lubbers-Khandekar, ***the configuration management*** (*e.g.*, system configuration by the control terminal 300 and its virtual machine control program 301) ***includes application routing for virtual assets or implementation of application rules*** (e.g., dynamic deployment or allocation of applications to virtual assets based on business policies and workload conditions).

XII. CONCLUSION

209. For all the reasons I have noted in the foregoing paragraphs, claims 1-20 of the '667 Patent are obvious in view of the references discussed above.

210. I currently hold the opinions set expressed in this declaration. But my analysis may continue, and I may acquire additional information and/or attain supplemental insights that may result in added observations.

APPENDIX A

Michael Mitzenmacher
michaelm@eecs.harvard.edu
33 Cary Avenue, Lexington MA 02421
617-496-7172

Research Interests Design and Analysis of Algorithms; Machine Learning Systems; Networks and Data Transmission; Databases.

Education

UNIVERSITY OF CALIFORNIA AT BERKELEY, Berkeley, CA
Ph.D. in Computer Science awarded December, 1996.
Dissertation: **The Power of Two Choices in Randomized Load Balancing.**
Advisor: Alistair Sinclair. GPA: 4.0/4.0

CAMBRIDGE UNIVERSITY, Cambridge, England
Attended as one of ten recipients of the Churchill Fellowship.
Cambridge C.A.S. in Mathematics with highest distinction awarded June 1992.

HARVARD COLLEGE, Cambridge, MA
B.A. in Mathematics with Computer Science, summa cum laude, awarded June 1991.

Employment

HARVARD UNIVERSITY, Cambridge, MA *Spring 1999-present*
Assistant professor (from Jan. 1999 -July 2002), **Associate professor** (from July 2002-January 2005), **Professor** (from Jan. 2005-present), **Area Dean for Computer Science** (from July 2010-June 2013), **Area Co-Chair for Computer Science** (from July 2018-June 2019).

Teach undergraduate courses including "Introduction to Algorithms and Data Structures" and "Discrete Mathematics for Computer Science", and graduate courses covering topics in randomized algorithms, algorithms for networks, compression, coding, cryptography, and information retrieval.

DIGITAL SYSTEMS RESEARCH CENTER, Palo Alto, CA *Fall 1996-Winter 1998*
Research scientist. Projects included work on information retrieval on the Web, erasure codes, error-correcting codes, on-line algorithms, and load balancing. Co-inventor for twelve submitted patents.

SANTA CLARA UNIVERSITY, Santa Clara, CA *Spring 1997*
Guest professor for the undergraduate class "Introduction to Algorithms."

Consultant: I consult on intellectual property issues as an expert witness and in other capacities. As an expert witness, I have testified in multiple trials. I have also consulted for several technology companies and research laboratories, including Adverplex (Cogolabs), Akamai, AT&T, Digital Fountain, eHarmony, Fluent Mobile (Fiksu), Google, Huawei, ITA Software, JobSync, Microsoft, Mitsubishi Research Laboratories, Weta Digital, and Yahoo.

Funding

NSF CCF-2101140: AF: Small: Algorithms and Data Structures with Predictions. PI: Michael Mitzenmacher. Total grant: $400,000. 7/21-6/24.

NSF CNS-2107078: CNS Core: Medium: Approximation and Randomization in the Programmable Data Plane. co-PIs: Minlan Yu. Total grant: $1,2000,000. 7/21-6/25.

NSF DMS-2023528: Foundations of Data Science Institute (Harvard University). 9/20-8/25.

NSF CCF-1563710: CIF: NeTS: Medium: Collaborative Research: Unifying Data Synchronization. co-PIs: David Starobinksi, Ari Trachtenberg. Total grant: $400,000. 7/16-6/20.

NSF CCF-1535705: AitF: FULL: Collaborative Research: Better Hashing for Applications: From Nuts & Bolts to Asymptotics. co-PIs: David Andersen. Total grant: $250,000. 9/15-8/19.

NSF CCF-1320231: AF: Small: Data Synchronization : Theory, Algorithms, and Practice PI: Michael Mitzenmacher. Total grant: $399,370. 9/13-8/16.

NSF CNS-1228598: TWC: Medium: Collaborative: Privacy-Preserving Distributed Storage and Computation. co-PIs: Michael Goodrich, Roberto Tamassia. Total grant: $400,000. 8/12-8/16.

NSF IIS-0964473: HCC: Medium: Collaborative Research: Data-Parallel Hash Tables: Theory, Practice and Applications. co-PIs: Nina Amenta. Total grant: $171,095. 8/10-7/13.

NSF CCF-0915922: AF : Small : The Theory and Practice of Hash-Based Algorithms and Data Structures. PI: Michael Mitzenmacher. Total grant: $441,956. 8/09-7/12.

NSF CNS-0721491: NeTS FIND: A Network-Wide Hashing Infrastructure for Monitoring and Measurement. PI: Michael Mitzenmacher. Total grant: $330,000. 9/07-8/10.

NSF CCF-0634923: Towards a Basic Understanding of Channels with Synchronization Errors. PI: Michael Mitzenmacher. Total grant: $200,000. 9/06-8/09.

NSF CCR-0121154: ITR/SY Algorithmic Issues in Large Scale Dynamic Networks. co-PIs: Eli Upfal, Brown. Subcontract to Harvard. Total grant: $502,000. 9/01-8/06.

NSF CCR-0118701: Low Density Parity Check Codes for Channels with Memory. co-PIs: Alek Kavcic. Total grant: $510,000. 9/01-8/04.

NSF CCR-9983832: Dynamic Processes and Network Algorithms (CAREER). $200,000. 7/00-6/04.

Dean's Competitive Fund for Promising Scholarship. 2018. $49,333.

Google University Research Program. 8/13-8/14. PIs: Eddie Kohler and Michael Mitzenmacher $56,000.

Google University Research Program. 12/11-12/12. $25,000.

Yahoo! University Research Program. 7/11-6/12. $10,000.

Google University Research Program. 12/09-12/10. $60,000.

Yahoo! University Research Program. 9/09-8/10. $10,000.

Google University Research Program. 8/08-7/09. $75,000.

Yahoo! University Research Program. 9/07-8/08. $25,000.

Yahoo! University Research Program. 9/06-8/07. $50,000.

Cisco University Research Program. 8/08-7/09. $80,000.

Cisco University Research Program. 8/07-7/08. $83,000.

Cisco University Research Program. 8/06-7/07. $80,000.

Cisco University Research Program. 8/05-7/06. $72,000.

Alfred P. Sloan Research Fellowship. $40,000. Awarded in 2000.

IBM Faculty Research Grant. $10,000 Awarded in 2005.

IBM Faculty Research Grant. $10,000 Awarded in 2003.

Mitsubishi Electronic Research Laboratory. $10,000 for undergraduate research projects. Awarded in 2002.

Honors

Applied Cryptography and Network Security Conference Test of Time Paper Award (2025)
ACM Paris Kanellakis Theory and Practice Award (2020)
IEEE Fellow (2020)
IEEE Senior Member (2018)
ACM Fellow (2014)
ACM Symposium on Parallelism in Algorithms and Architectures Best Paper Award (2014)
World Wide Web Conference Best Paper Award (2014)
Royal Academy of Engineering Distinguished Visiting Fellowship (2010)
ACM SIGCOMM Test of Time Paper Award (2009)
IEEE Information Theory Society Best Paper Award (2002)
Alfred P. Sloan Research Fellowship (2000)
Sakrison Award (for Ph.D. thesis at Berkeley) (1997)
NDSEG Graduate Fellowship (1992-95)
NSF Graduate Fellowship Winner (1991)
Churchill Fellowship (1991-92)
Hoopes Prize (for senior thesis at Harvard) (1991)
Phi Beta Kappa (1990)
Harvard Distinction in Teaching Award (1990)
Goldwater Fellowship (1989-91)

Professional Activities

Editorships and Boards:
Advisory Board, Hariri Institute at Boston University (2018-2021)
Editorial Board, Leibniz International Proceedings in Informatics (2009-2019)
Editorial Board, Communications of the ACM (2013-2018)
Science Board, Santa Fe Institute (2013-2019)
Guest Editor, SIAM Journal of Computing special issue for STOC 2009
SIAM Journal on Computing, Editor (2006-2010)
Internet Mathematics, Managing Editor
Guest Editor, Theory of Computing System special issue for SPAA 2002
Journal of Interconnection Networks, Editor

Organizing Committees:

Sublinear Algorithms Workshop Organizing Committee, FODSI (2022)
Algorithms with Predictions Workshop Organizing Committee, STOC (2022)
Machine Learning for Algorithms Workshop Organizing Committee, FODSI (2021)
SIGACT Chair (2015-2018)
Dagstuhl Seminar: Theory and Applications of Hashing (2017)
EADS Summer School on Hashing: Theory and Applications (2014)
ICERM Workshop on Stochastic Graph Models (2014)
WAW 2013 Organizing Committee
SIGACT Vice-Chair (2009-2012)
General Chair, STOC (2010)
Workshop on Randomized Algorithms and Random Graphs (2010)
DARPA Information Science and Technology Study Groups (2007-2008)
SIGACT Committee for the Advancement of Theoretical Computer Science (2005-2008)
MSRI Workshop on Models of Real-World Random Networks (2005)
FOCS 2003 Local Arrangements Chair
Second Workshop on Randomized Algorithms and Random Graphs (2003)
Workshop on Algorithms and Models for the Web Graph (2002)
First Workshop on Randomized Algorithms and Random Graphs (2002)
DIMACS Workshop on Quality of Service Issues in the Internet (2001)
BU/NSF Workshop on Internet Measurement, Instrumentation, and Characterization (1999)
FOCS 1998 Local Arrangements Chair

Program Committees:
ACM CoNEXT 2026 Program Committee
ICLR 2026 Program Committee
NSDI 2025 Program Committee
HotNets 2024 Program Committee
WADS 2023 Program Committee
ACM CoNEXT 2023 Program Committee
NeurIPS 2022 Program Committee
ACM CoNEXT 2022 Program Committee
ICML 2022 Program Committee
SIGCOMM 2022 Program Committee
NeurIPS 2021 Program Committee
ICML 2021 Program Committee
ACDA 2021 Program Committee
ICLR 2021 Program Committee
APoCS 2021 Program Committee
ACM CoNEXT 2020 Program Committee
NeurIPS 2020 Program Committee
ICML 2020 Program Committee
APoCS 2020 Program Committee
NeurIPS 2019 Program Committee
SIGCOMM 2019 Program Committee
SOSA 2019 Program Committee (Co-Chair)
ACM CoNEXT 2018 Program Committee
ANALCO 2018 Program Committee
SIGCOMM 2017 Program Committee
NSDI 2017 Program Committee
ACM CoNEXT 2016 Program Committee
SIGCOMM 2016 Program Committee

ALENEX 2016 Program Committee (Co-Chair)
ICALP 2016 Program Committee (Track C PC Chair)
ACM CoNEXT 2015 Program Committee
ICALP 2015 Program Committee
ITCS 2015 Program Committee
SIGCOMM 2014 Program Committee
ACM CoNEXT 2014 Program Committee
WSDM 2014 Senior Program Committee
WAW 2013 Program Committee (Co-Chair)
STOC 2013 Executive Committee
WSDM 2013 Senior Program Committee
SIGCOMM 2012 Program Committee
ISIT 2012 Program Committee
NSDI 2012 Program Committee
WAW 2011 Program Committee
ACM CoNEXT 2010 Program Committee
SIGCOMM 2010 Program Committee
NSDI 2010 Program Committee
FUN 2010 Program Committee
LATIN 2010 Program Committee
NetSciCom 2009 Program Committee
SIGCOMM 2009 Program Committee
STOC 2009 Program Committee (PC Chair)
NSDI 2009 Program Committee
ICALP 2008 Program Committee
SPAA 2008 Program Committee
WAW 2007 Program Committee
Analysis of Algorithms 2007 Program Committee
ANALCO 2007 Program Committee
ICALP 2007 Program Committee
RANDOM 2007 Program Committee
NCA 2006 Program Committee
STOC 2006 Program Committee
PODC 2005 Program Committee
ANALCO 2005 Program Committee
PODC 2004 Program Committee
SIGCOMM 2004 Program Committee
IPTPS 2004 Program Committee
Data Compression Conference 2004 Program Committee
ESA 2004 Program Committee
WAW 2003 Program Committee
FOCS 2003 Program Committee
Data Compression Conference 2003 Program Committee
ESA 2002 Program Committee
PODC 2002 Program Committee
SPAA 2002 Program Committee
STOC 2002 Program Committee
Data Compression Conference 2002 Program Committee
ALENEX 2002 Program Committee
HiPC 2001 Program Committee
ISAAC 2000 Program Committee

RANDOM 2000 Program Committee
STACS 2000 Program Committee
FOCS 99 Program Committee
RANDOM 98 Program Committee

Reviewer for several conferences, journals, and grant panels.

Books

Probability and Computing: Randomized Algorithms and Probabilistic Analysis, by Michael Mitzenmacher and Eli Upfal. Published by Cambridge University Press in 2005. (2nd edition available in 2017.) This is a textbook meant for an advanced undergraduate or beginning graduate class. The textbook has been used in courses at Brown, Harvard, U. C. Berkeley, Univ. of Victoria, Tufts, Univ. of Mass. at Amherst, Purdue, U. Penn., and several other universities.

Conference and Journal Publications

"The Mixed Birth-death/death-Birth Moran Process," with D. Brewster, Y. Huang, and M. Nowak. To appear in *Proceedings of 17th Innovations in Theoretical Computer Science Conference*, 2026.

"Spatiotemporal Sketch Disaggregation: Streaming Analytics with Heterogeneous Resources," with J. Langlet, P. Chen, Z. Liu, R. Basat, and G. Antichi. To appear in *Proceedings of the 42nd IEEE International Conference on Data Engineering*, 2026.

"Better than Optimal: Improving Adaptive Stochastic Quantization Using Shared Randomness," with R. Basat, Y. Ben-Itzhak, and S. Vargaftik. To appear in *Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS)*, 2026.

"Call on MARS: Scheduling API-Augmented LLM Requests," with R. Shahout, C. Liang, S. Xin, Q. Lao, Y. Cui, and M. Yu. To appear in *Proceedings of the 39th Conference on Neural Information Processing Systems*, 2025.

"HACK: Homomorphic Acceleration via Compression of the Key-Value Cache for Disaggregated LLM Inference," with Z. Zhang, H. Shen, S. Vargarftik, R. Basat, and M. Yu. To appear in *Proceedings of ACM SIGCOMM*, 2025.

"Queueing, Predictions, and Large Language Models: Challenges and Open Problems," with R. Shahout. To appear in *Stochastic Systems*.

"Don't Stop Me Now: Embedding Based Scheduling For LLMs," with R. Shahout, E. Malach, C. Liu, W. Jiang, and M. Yu. In *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025.

"Parallel Peeling of Invertible Bloom Lookup Tables in a Constant Number of Rounds," with M. Goodrich and R. Kitagawa. In *Proceedings of the 50th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, pp. 70-84, 2025.

"Efficient d-ary Cuckoo Hashing at High Load Factors by Bubbling Up," with W. Kuszmaul. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pp. 3931-3952, 2025.

"Leveraging Parameterized Chernoff Bounds for Simplified Algorithm Analyses," with M. Dillencourt and M. Goodrich. *Information Processing Letters*, 187:106516, 2025.

"Optimal and Approximate Adaptive Stochastic Quantization," with R. Basat, Y. Ben-Itzhak, and S. Vargaftik. In *Proceedings of the 38th Conference on Neural Information Processing Systems*, 2024.

"SkipPredict: When to Invest in Predictions for Scheduling," with R. Shahout. In *Proceedings of the 38th Conference on Neural Information Processing Systems*, 2024.

"Beyond Throughput and Compression Ratios: Towards High End-to-end Utility of Gradient Compression," with W. Han, S. Vargaftik, B. Karp, and R. Basat. In *Proceedings of the Twenty-Third ACM Workshop on Hot Topics in Networks*, pp. 186-194, 2024.

"Learning-Augmented Frequency Estimation in Sliding Windows," with R. Shahout and I. Sabek. In *Proceedings of the Workshop on Intelligent Classification of High-Speed Network Traffic (at ICNP '24).*

"Learning-Based Heavy Hitters and Flow Frequency Estimation in Streams," with R. Shahout. In *Proceedings of the 32nd IEEE International Conference on Network Protocols*, 2024.

"Tabula: Efficiently Computing Nonlinear Activation Functions for Secure Neural Network Inference," with M. Lam, V. J. Reddi, G. Wei, and D. Brooks. In *Transactions on Machine Learning Research*, 2024.

"Accelerating Federated Learning with Quick Distributed Mean Estimation," with S. Vargaftik, R. Basat, A. Portnoy, G. Mendelson, and Y. Ben-Itzhak. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 3410-3442, 2024.

"THC: Accelerating Distributed Deep Learning Using Tensor Homomorphic Compression," with M. Li, R. Basat, S. Vargaftik, C. Lao, K. Xu, X. Tang, and, M. Yu. In *Proceedings of NSDI*, pp. 1191-1211, 2024.

"Separating k-Player from t-Player One-Way Communication, with Applications to Data Streams," with E. Du, D. Woodruff, and G. Yang. In *Theory of Computing*, vol. 19, Article 10, pp. 1-44, 2023.

"Direct Telemetry Access," with J. Langlet, R. Basat, G. Oliaro, M. Yu, and G. Antichi. In *Proceedings of ACM SIGCOMM*, pp. 832-849, 2023.

"Can Learned Models Replace Hash Functions?" with I. Sabek, K. Vaidya, D. Horn, A. Kipf, and T. Kraska. *Proceedings of the VLDB Endowment*, vol. 16, no. 3, pp. 532-545, 2022.

"Algorithmic Tools for Mining the Motif Structure of Networks," with B. Matejek, T. Chen, and C. Tsourakakis. To appear in *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2022.

"Edge-colored directed subgraph enumeration on the connectome," with B. Matejek, D. Wei, T. Chen, C. Tsourakakis, and H. Pfister. *Scientfic Reports*, 2022.

"Algorithms with Predictions," with S. Vassilvitskii. *Communications of the ACM*, Volume 65 Issue 7, pp. 33-35, 2022.

"EDEN: Communication-Efficient and Robust Distributed Mean Estimation for Federated Learning," with S. Vargaftik, R. Basat, A. Portnoy, G. Mendelson, and Y. Ben-Itzhak. In

Proceedings of the 39th International Conference on Machine Learning, pp. 21984-22014, 2022.

"SNARF: A Learning-Enhanced Range Filter," with K. Vaidya, S. Chatterjee, E. Knorr, S. Idreos, and T. Kraska. In *Proceedings of the VLDB Endowment*, 15(8):1632-1664, 2022.

"Proteus: A Self-Designing Range Filter," with E. Knorr, B. Lemaire, A. Lim, S. Luo, H. Zhang, and S. Idreos. In *Proceedings of the SIGMOD International Conference on Management of Data*, pp. 1670-1684, 2022.

"Designing Sensitive Viral Diagnostics with Machine Learning," with H. Metsky, N. Welch, P. Pillai, N. Haradhvala, L. Rumker, S. Mantena, Y. Zhang, D. Yang, C. Ackerman, J. Weller, P. Blainey, C. Myhrvold, and P. Sabeti. *Nature Biotechnology*, 2022.

"The Supermarket Model with Known and Predicted Service Times," with M. Dell'Amico. *IEEE Transactions on Parallel and Distributed Systems*, 2022.

"Uniform Bounds for Scheduling with Job Size Estimates," with Z. Scully and I. Grosof. In *Proceedings of 13th Innovations in Theoretical Computer Science Conference*, 114:1-114:30, 2022.

"DRIVE: One-bit Distributed Mean Estimation," with S. Vargaftik, R. Basat, A. Portnoy, G. Mendelson, and Y. Ben-Itzhak. In *Proceedings of the 34th Conference on Neural Information Processing Systems*, pp. 362-377, 2021.

"Zero-CPU Collection with Direct Telemetry Access," with J. Langlet, R. Basat, S.S. Ramanathan, G. Oliaro, M. Yu, and G. Antichi. In *Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks*, pp. 108-115, 2021.

"Gradient Disaggregation: Breaking Privacy in Federated Learning by Reconstructing the User Participant Matrix," with M. Lam, V. J. Reddi, D. Brooks, and G. Wei. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 5959-5968, 2021.

"Putting the 'Learning' into Learning-Augmented Algorithms for Frequency Estimation," with E. Du and F. Wang. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 2860-2869, 2021.

"Queues with Small Advice." In *Proceedings of the SIAM Conference on Applied and Computational Discrete Algorithms*, pp. 1-12, 2021.

"How to Send a Real Number Using a Single Bit (and Some Shared Randomness)," with R. Basat and S. Vargaftik. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming*, 25:1-25:20, 2021.

"SALSA: Self-Adjusting Lean Streaming Analytics," with R. Basat, G. Einzinger, and S. Vargaftik. In *The 37th IEEE International Conference on Data Engineering*, pp. 864-875, 2021.

"Partitioned Learned Bloom Filters," with K. Vaidya, E. Knorr, and T. Kraska. In *The Ninth International Conference on Learning Representations (ICLR)*, 2021.

"Improved Sublinear Time Algorithm for Longest Increasing Subsequence," with S. Seddighin. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pp. 1934-1947, 2021.

"DISCOvering the heavy hitters with disaggregated sketches," with V. Bruschi, R. Basat, Z. Liu, G. Antichi, and G. Bianchi. In *Proceedings of the 16th International Conference on Emerging Networks Experiments and Technologies (CoNEXT)*, pp. 536-537, 2020.

"Detecting Routing Loops in the Data Plane," with J. Kucera, R. Basat, M. Kuka, G. Antichi, and M. Yu. In *Proceedings of the 16th International Conference on Emerging Networks Experiments and Technologies (CoNEXT)*, pp. 466-473, 2020.

"Optimal Learning of Joint Alignments with a Faulty Oracle," with K. Larsen and C. Tsourakakis. In *Proceedings of ISIT*, pp.2492-2497, 2020.

"PINT: Probabilistic In-band Network Telemetry," with R. Basat, S. Ramanathan, Y. Li, G. Antichi, and M. Yu. In *Proceedings of ACM SIGCOMM*, pp. 662-680, 2020.

"Dynamic Algorithms for LIS and Distance to Monotonicity," with S. Seddighin. In *Proccedings of the 52nd Annual Symposium on Theory of Computing*, pp. 671-684, 2020.

"Clustering with a Faulty Oracle," with K. Larsen and C. Tsourakakis. In *Proceedings of the Web Conference*, pp. 2831-2834, 2020.

"Prophets, Secretaries, and Maximizing the Probability of Choosing the Best," with H. Esfandiari, M. Hajiaghayi, and B. Lucier. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pp. 3717-3727, 2020.

"Faster and More Accurate Measurement through Additive-Error Counters," with R. Basat, G. Einziger, and S. Vargaftik. In *Proceedings of IEEE INFOCOM*, pp. 1251-1260, 2020. o "Scheduling with Predictions and the Price of Misprediction." In *Proceedings of 11th Innovations in Theoretical Computer Science Conference*, 14:1-14:18, 2020.

"Equitability, Interval Estimation, and Statistical Power," with Y. Reshef, D. Reshef, and P. Sabeti. *Statistical Science*, (35)2:202-217, 2020.

"A Digital Fountain Retrospective," with J. Byers and M. Luby. *ACM SIGCOMM Computer Communication Review*, 49(5), pp. 82-85, 2019.

"Robust Set Reconciliation via Locality Sensitive Hashing," with T. Morgan. In *Proceedings of the 38th Symposium on Principles of Database Systems*, pp. 164-181, 2019.

"Online Pandora's Boxes and Bandits," with H. Esfandiari, M. Hajiaghayi, and B. Lucier. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pp. 1885-1892, 2019.

"Arithmetic Progression Hypergraphs: Examining the Second Moment Method." In *Proceedings of ANALCO*, pp. 127-134, 2019.

"A Bayesian Nonparametric View on Count-Min Sketch," with R. Adams and D. Cai. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, pp. 8782-8791, 2018.

"A Model for Learned Bloom Filters and Optimizing by Sandwiching." In *Proceedings of the 31st Conference on Neural Information Processing Systems*, pp. 462-471, 2018.

"Metric Sublinear Algorithms via Linear Sampling," with H. Esfandiari. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*, pp. 11-22, 2018.

"EMOMA: Exact Match in One Memory Access," with S. Pontarelli and P. Reviriego. *IEEE Transactions on Knowledge and Data Engineering*, 30(11):2120-2133, 2018.

"Weightless: Lossy Weight Encoding for Deep Neural Network Compression," with B. Reagen, U. Gupta, R. Adolf, A. Rush, G. Wei, and D. Brooks. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 4321-4330, 2018.

"Load Thresholds for Cuckoo Hashing with Double Hashing," with K. Panagiotou and S. Walzer. In *Proceedings of the 16th Scandinavian Symposium and Workshops on Algorithm Theory*, pp. 29:1-29:9, 2018.

"Joint Alignment From Pairwise Differences with a Noisy Oracle," with C. Tsourakakis. In *Proceedings of the 15th Workshop on Algorithms and Models for the Web Graph*, pp. 59-69, 2018. Journal version: *Internet Mathematics*, at https://doi.org/10.24166/im.06.2019, December 2020.

"Reconciling Graphs and Sets of Sets," with T. Morgan. In *Proceedings of the 37th Symposium on Principles of Database Systems*, pp. 33-47, 2018.

"Simulated Annealing for JPEG Quantization," with M. Hopkins and S. Wagner-Carena. In *Proceedings of the 2018 IEEE Data Compression Conference*, p. 414.

"Adaptive Cuckoo Filters," with S. Pontarelli and P. Reviriego. In *Proceedings of the ALENEX Conference*, pp. 36-47, 2018. Journal version: *ACM Journal of Experimental Algorithmics*, Article 1.1, March 2020.

"An Empirical Comparison of the Maximal and Total Information Coefficients to Leading Measures of Dependence," with D. Reshef, Y. Reshef, and P. Sabeti. *Annals of Applied Statistics*, 12(1):123-155, 2018.

"Simple Multi-Party Set Reconciliation," with R. Pagh. *Distributed Computing* 31(6):441-453, 2018.

"Technical Perspective: Building a Better Hash Function." *Communications of the ACM*, 60(7), p. 93, 2017.

"Compresso: Efficient Compression of Segmentation Data For Connectomics," with D. Haehn, F. Lekschas, B. Matejek, and H. Pfister. In *Proceedings of the 20th International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 781-788, 2017.

"Auditable Data Structures," with M. Goodrich, E. Kornaropoulos, and R. Tamassia. In *Proceedings of the European Symposium on Security and Privacy*, pp. 285-300, 2017.

"2-3 Cuckoo Filters for Faster Triangle Listing and Set Intersection," with D. Eppstein, M. Goodrich, and M. Torres. In *Proceedings of the 36th Symposium on Principles of Database Systems*, pp. 247-260, 2017.

"Scalable motif-aware graph clustering," with C. Tsourakakis and J. Pachocki. In *Proceedings of the 26th International Conference on the World Wide Web*, pp. 1451-1460, 2017.

"Measuring Dependence Powerfully and Equitably," with Y. Reshef, D. Reshef, H. Finucane, and P. Sabeti. *Journal of Machine Learning Research*, 17(212):1-63, 2016.

"Quantized Random Projections and Non-Linear Estimation of Cosine Similarity," with P. Li and M. Slawski. In *Advances In Neural Information Processing Systems 29*, pp.2748-2756, 2016.

"Analyzing Distributed Join-Idle-Queue: A Fluid Limit Approach." In *Proceedings of the 54th Annual Allerton Conference on Communication, Control, and Computing*, pp. 312-318, 2016.

"Hardness of Peeling with Stashes," with V. Nathan. *Information Processing Letters*, 116(11), pp. 682-688, 2016.

"More Practical and Secure History-Independent Hash Tables," with M. Goodrich, E. Kornaropoulos, and R. Tamassia. In *European Symposium on Research in Computer Security*, pp. 20-38, 2016.

"Models and Algorithms for Graph Watermarking," with D. Eppstein, M. Goodrich, J. Lam, N. Mamano, and M. Torres. In *Proceedings of the Information Security Conference*, pp. 283-301, 2016.

"Better Bounds for Coalescing-Branching Random Walks on Graphs," with R. Rajaraman and S. Roche. In *Proceedings of the 28th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 313-323, 2016. Journal version: *ACM Transactions on Parallel Computing*, 5(1), 2:1-2:23, 2018.

"OMASS: One Memory Access Set Separation," with S. Pontarelli and P. Reviriego. *IEEE Transactions on Knowledge and Data Engineering*, 28(7):1940-1943, 2016.

"Voronoi Choice Games", with M. Boppana, R. Hod, and T. Morgan. In *Proceedings of 43rd International Colloquium on Automata, Languages, and Programming*, 23:1-23:13, 2016.

"Space Lower Bounds for Itemset Frequency Sketches," with E. Liberty, J. Thaler, and J. Ullman. In *Proceedings of the 35th Symposium on Principles of Database Systems*, pp. 441-454, 2016.

"Technical Perspective: Catching Lies (and Mistakes) in Offloaded Computation," with J. Thaler. *Communications of the ACM*, 59(2), p. 102, 2016.

"A New Approach to Analyzing Robin Hood Hashing." In *Proceedings of ANALCO*, pp. 10-24, 2016.

"More Analysis of Double Hashing for Balanced Allocations." In *Proceedings of ANALCO*, pp. 1-9, 2016.

"Theory Without Experiments: Have We Gone Too Far?" *Communications of the ACM*, 58(9), pp. 40-42, 2015.

"Scaling Up Clustered Network Appliances with XBricks," with D. Zhou, B. Fan, H. Lim, R. Wang, M. Kaminsky, D. Andersen, and A. Singh. In *Proceedings of ACM SIGCOMM*, 2015.

"Scalable Large Near-Clique Detection in Large-Scale Networks via Sampling," with J. Pachocki, R. Peng, C. Tsourakakis, and S. Chen Xu. In *Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2015.

"Repeated Deletion Channels," with B. Haeupler. In *Proceedings of the IEEE Information Theory Workshop*, pp. 152-156, 2014.

"Multi-Party Set Reconciliation Using Characteristic Polynomials," with A. Boral. In *Proceedings of the 52nd Annual Allerton Conference on Communication, Control, and Computing*, pp. 1182-1187, 2014.

"Cleaning Up the Record on the Maximal Information Coefficent and Equitability," with D. Reshef, Y. Reshef, and P. Sabeti. *Proceedings of the National Academy of Sciences* 111(33):E3362-3, 2014.

"Cuckoo Filter: Practically Better than Bloom," with B. Fan, D.G. Andersen, and M. Kaminsky. In *Proceedings of the 10th International Conference on Emerging Networks Experiments and Technologies (CoNEXT)*, pp. 75-88, 2014.

"Coding for Random Projections," with P. Li and A. Shrivastava. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pp. 676–684, 2014.

"Balanced Allocations and Double Hashing." In *Proceedings of the 26th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 331-342, 2014.

"Parallel Peeling Algorithms," with J. Jiang and J. Thaler. In *Proceedings of the 26th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 319-330, 2014. Journal version: *ACM Transactions on Parallel Computing*, 3(1), 7, 2016.

"Wear Minimization for Cuckoo Hashing: How Not to Throw a Lot of Eggs into One Basket", with D. Eppstein, M. Goodrich, and P. Pszona. In *Proceedings of the 13th International Symposium on Experimental Algorithms*, pp. 162-173, 2014.

"Efficient Estimation for High Similarities using Odd Sketches," with R. Pagh and N. Pham. In *Proceedings of the 23rd International Conference on the World Wide Web*, pp. 109-118, 2014.

"Improving the Performance of Invertible Bloom Lookup Tables," with S. Pontarelli and P. Reviriego. *Information Processing Letters*, 114(4), pp. 185-191, 2014.

"Special Section on the Forty-First Annual ACM Symposium on Theory of Computing (STOC 2009)," with N. Immorlica, J. Katz, R. Servedio, and C. Umans. *SIAM Journal on Computing*, 41(6):1591-1592, 2012.

"The Daily Deals Marketplace: Empirical Observations and Managerial Implications," with J. Byers and G. Zervas. *AC SIGecom Exchanges*, 11(2):29-31, 2012.

"Peeling Arguments and Double Hashing," with J. Thaler. In *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1118-1125, 2012.

"The Complexity of Object Reconciliation, and Open Problems Related to Set Difference and Coding," with G. Varghese. In *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1126-1132, 2012.

"Cache-Oblivious Dictionaries and Multimaps with Negligible Failure Probability," with M. Goodrich, D. Hirschberg, and J. Thaler. In *Proceedings of the First Mediterranean Conference on Algorithms*, pp. 203-218, 2012.

"An Economic Analysis of User-Privacy Options in Ad-Supported Services," with J. Feigenbaum and G. Zervas. In *Proceedings of the 8th Workshop on Internet and Network Economics*, pp. 30-43, 2012.

"Verifiable Computation with Massively Parallel Interactive Proofs," with H. Pfister, M. Roberts, and J. Thaler. In *Proceedings of the USENIX HotCloud Workshop, 2012.*

"Biff (Bloom Filter) Codes: Fast Error Correction for Large Data Sets," with G. Varghese. In *Proceedings of the International Symposium on Information Theory*, pp. 483-487, 2012.

"Continuous Time Channels with Interference," with I. Ivan, J. Thaler, and H. Yuen. In *Proceedings of the International Symposium on Information Theory*, pp. 860-864, 2012.

"The Groupon Effect on Yelp Ratings: a Root Cause Analysis," with J. Byers and G. Zervas. In *Proceedings of the ACM Conference on Electronic Commerce*, pp. 248-265, 2012.

"Anonymous Card Shuffling and Its Applications to Parallel Mixnets," with M. Goodrich. In *Proceedings of 39th International Colloquium on Automata, Languages, and Programming*, pp. 549-560, 2012.

"Chernoff-Hoeffding Bounds for Markov Chains: Generalized and Simplified," with K. Chung, H. Lam, and Z. Liu. In *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science*, pp. 124-135, 2012.

"Practical oblivious storage," with M. Goodrich, O. Ohrimenko and R. Tamassia. In *Proceedings of the Second ACM Conference on Data and Application Security and Privacy*, pp. 13-24, 2012.

"Hierarchical Heavy Hitters with the Space Saving Algorithm," T. Steinke and J. Thaler. In *Proceedings of 2012 Meeting on Algorithm Engineering and Experiments*, pp. 160-174.

"Daily Deals: Prediction, Social Diffusion, and Reputational Ramifications," with J. Byers and G. Zervas. In *Proceedings of the 5th International Conference on Web Search and Web Data Mining*, pp. 543-552, 2012.

"Practical Verified Computation with Streaming Interactive Proofs," with G. Cormode and J. Thaler. In *Proceedings of Innovations in Theoretical Computer Science*, pp. 90-112, 2012.

"Information Dissemination via Random Walks in d-Dimensional Space," with H. Lam, Z. Liu, X. Sun, and Y. Wang. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1612-1622, 2012.

"Privacy-Preserving Group Data Access via Stateless Oblivious RAM Simulation," with M.Goodrich and R. Tamassia. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 157-167, 2012.

"An Efficient Rigorous Approach for Identifying Statistically Significant Frequent Itemsets," with A. Kirsch, A. Pietracaprina, G. Pucci, E. Upfal, and F. Vandin. *Journal of the ACM*, 59(3), 2012.

"Detecting Novel Associations in Large Datasets," with D. Reshef, Y. Reshef, H. Finucane, S. Grossman, G. McVean, P. Turnbaugh, E. Lander, and P. Sabeti. *Science*, Vol. 3004, No. 6602, pp. 1518-1524, December 2011.

"External-Memory Multimaps," with E. Angelino, M. Goodrich, and J. Thaler. In *Proceedings of the 22nd International Symposium on Algorithms and Computation*, pp. 384-394, 2011. Journal version: *Algorithmica*, 67(1), pp. 23-48, 2013.

"Brief announcement: large-scale multimaps," with M. Goodrich. In *Proceedings of the 23rd ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 259-260, 2011.

"Oblivious RAM Simulation with Efficient Worst-Case Access Overhead," with M. Goodrich, O. Ohrimenko, and R. Tamassia. In *Proceedings of 3rd ACM Cloud Computing Workshop*, pp. 95-100.

"Invertible Bloom Lookup Tables," with M. Goodrich. In *Proceedings of the 49th Annual Allerton Conference on Communication, Control, and Computing*, pp. 792-799, 2011.

"Cuckoo Hashing with Pages," with M. Dietzfelbinger and M. Rink. In *Proceedings of the 19th Annual Europeans Symposium on Algorithms*, pp. 615-627, 2011.

"Network Coding Meets TCP: Theory and Implementation," with J.K. Sundararajan, D. Shah, M. Medard, S. Jakubczak, and J. Barros. *Proceedings of the IEEE*, 99(3), pp. 490-510, 2011.

"Privacy-Preserving Access of Outsourced Data via Oblivious RAM Simulation," with M. Goodrich. In *Proceedings of 38th International Colloquium on Automata, Languages, and Programming*, pp. 576-587, 2011.

"Graption: A Graph-based P2P Traffic Classification Framework for the Internet Backbone," with M. Iliofotou, H. Kim, M. Faloutsos, P. Pappu, and G. Varghese. *Computer Networks*, 55(8), pp. 1909-1920, 2011.

"On the Zero-Error Capacity Threshold for Deletion Channels," with I. Kash, J. Thaler, J. Ullman. In *Proceedings of the Information Theory and Applications Workshop*, pp. 1-5, 2011.

"Heapable Sequences and Subsequences," with J. Byers, B. Heeringa, and G. Zervas. In *Proceedings of ANALCO, pp. 33-44, 2011.*

"An Introduction to Human Guided Search." *ACM Crossroads*, 17(2):34-35, 2010.

"Popularity is Everything: A New Approach to Protecting Passwords from Statistical Guessing Attacks," with S. Schechter and C. Herley. In *Proceedings of the 5th Usenix Workshop on Hot Topics in Security*, pp. 1-8, 2010.

"Streaming Graph Computations with a Helpful Advisor," with G. Cormode and J. Thaler. In *Proceedings of the European Symposium on Algorithms*, pp. 231-242, 2010. Journal version: *Algorithmica*, 65(2), pp. 409-442, 2013.

"Tight Thresholds for Cuckoo Hashing via XORSAT," with M. Dietzfelbinger, A. Goerdt, A. Montanari, R. Pagh, and M. Rink. In *Proceedings of 37th International Colloquium on Automata, Languages, and Programming*, pp. 213-225, 2010.

"Local cluster aggregation models of explosive percolation," with R.M. D'Souza. *Physical Review Letters*, 14(19), p. 104-107, 2010.

"Tight Asymptotic Bounds for the Deletion Channel with Small Deletion Probabilities," with A. Kalai and M. Sudan. In *Proceedings of the International Symposium on Information Theory*, pp. 997-1001, 2010.

"An Improved Analysis of the Lossy Difference Aggregator," with H. Finucane. *Computer Communication Review*, 40(2), pp. 4-11, 2010.

"Information Asymmetries in Pay-Per-Bid Auctions: How Swoopo Makes Bank," with J. Byers and G. Zervas. In *Proceedings of the ACM Conference on Electronic Commerce*, pp. 1-12, 2010.

"Carousel: Scalable Logging for Intrusion Prevention Systems", with T. Lam and G. Varghese. In *Proceedings of NSDI 2010*, pp. 361-376.

"AMS Without 4-Wise Independence on Product Domains," with V. Braverman, K. Chung, Z. Liu, and R. Ostrovsky. In *Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science*, pp. 119-130, 2010.

"Adaptive Weighing Designs for Keyword Value Computation," with J. Byers and G. Zervas. In *Proceedings of the Third International Conference on Web Search and Web Data Mining* pp. 331-340, 2010.

"Human-guided search," with G. Klau, N. Lesh, and J. Marks. *Journal of Heuristics*, 16(3), pp. 289-310, 2010.

"Real-Time Parallel Hashing on the GPU," with D. Alcantara, A. Sharf, F. Abbasinejad, S. Sengupta, J. Owens, and N. Amenta. *ACM Transactions on Graphics*, 28(5), p. 154, 2009.

"Exploiting Dynamicity in Graph-based Traffic Analysis: Techniques and Applications," with M. Iliofotou and M. Faloutsos. In *Proceedings of ACM Co-NEXT*, pp. 241-252, 2009.

"An Analysis of Random-Walk Cuckoo Hashing," with A. Frieze and P. Mellsted. In *Proceedings of 2009 APPROX-RANDOM*, pp. 490-503, 2009. Journal version: *SIAM Journal on Computing*, 40(2), pp. 291-308, 2011.

"Some Open Questions Related to Cuckoo Hashing." In *Proceedings of the 17th Annual European Symposium on Algorithms*, pp. 1-10, 2009.

"On Compressing Social Networks,", with F. Chierichetti, R. Kumar, S. Lattanzi, A. Panconesi, and P. Raghavan. In *Proceedings of the 15th ACM SIGKDD Int'l Conference on Knowledge Discovery and Data Mining*, pp. 219-228, 2009.

"An Efficient Rigorous Approach for Identifying Statistically Significant Frequent Itemsets," with A. Kirsch, A. Pietracaprina, G. Pucci, E. Upfal, and F. Vandin. In *Proceedings of the 28th Symposium on Principles of Database Systems*, pp. 117-126, 2009.

"Using the Power of Two Choices to Improve Bloom Filters," with Steven Lumetta. *Internet Mathematics*, vol. 4. no. 1, pp. 17-33, 2009.

"Network coding meets TCP," with J. Sundararajan, D. Shah, M. Medard, and J. Barros. In *Proceedings of IEEE INFOCOM*, pp. 280-288, 2009.

"An Economically Principled Generative Model of AS Graph Connectivity," with J. Corbo, S. Jain, and D. Parkes. In *Proceedings of IEEE INFOCOM (Mini-conference), pp. 2941-2945, 2009.*

"Designing floating codes for expected performance," with H. Finucane and Z. Liu. In *Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computing*, 2008. Journal version (with F. Chierichetti, H. Finucane, and Z. Liu): *IEEE Transactions on Information Theory*, vol 56, Issue 3, pp. 968-978, 2010.

"On the performance of multiple choice hash tables with moves on deletions and inserts," with A. Kirsch. In *Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1284-1290, 2008.

"More Robust Hashing: Cuckoo Hashing with a Stash," with A. Kirsch and U. Wieder. In *Proceedings of the 16th Annual European Symposium on Algorithms*, pp. 611-622, 2008. Journal version: *SIAM Journal on Computing*, 39(4), pp. 1543-1561, 2009.

"A survey of results for deletion channels and related synchronization channels." In *Proceedings of the 2008 Scandinavian Workshop on Algorithm Theory*, pp. 1-3, 2008. Journal version: "New Results and Open Problem for Channels with Synchronization." *Probability Surveys*, 2009.

"The Power of One Move: Hashing Schemes for Hardware," with A. Kirsch. In *Proceedings of IEEE INFOCOM*, pp. 106-110, 2008. Journal version: *IEEE/ACM Transactions on Networking*, 18(6):1752-1765, 2010.

"Distributed Beamforming with Binary Signaling," with M. Johnson and K. Ramchandran. In *Proceedings of the 2008 Int'l Symposium on Information Theory (ISIT)*, pp. 890-894, 2008.

"The Hiring Problem and Lake Wobegon Strategies," with A. Broder, A. Kirsch, R. Kumar, E. Upfal, and S. Vassilvitskii. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1184-1193, 2008. Journal version: *SIAM Journal on Computing*, 39(4), pp. 1233-1255, 2009.

"Why Simple Hash Functions Work: Exploiting the Entropy in a Data Stream," with S. Vadhan. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 746-755, 2008. Journal version: with K. Chung and S. Vadhan. *Theory of Computing*, 9, pp. 897-945-1255, 2013.

"Trace reconstruction with constant deletion probability and related results," with T. Holenstein, R. Panigrahy, and U. Wieder. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 389-398, 2008.

"Capacity bounds for sticky channels." *IEEE Transactions on Information Theory*, vol 54, Issue 1, pp. 72-77, 2008.

"Network Monitoring Using Traffic Dispersion Graphs," with M. Iliofotou, P. Pappu, G. Varghese, M. Faloutsos, and S. Singh. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, pp. 315-320, 2007.

"HEXA: Compact Data Structures for Faster Packet Processing," with S. Kumar, J. Turner, and P. Crowley. In *Proceedings of the IEEE International Conference on Network Protocols*, pp. 246-255, 2007.

"Using a Queue to De-amortize Cuckoo Hashing in Hardware," with A. Kirsch and M. Mitzenmacher. In *Proceedings of the 45th Annual Allerton Conference on Communication, Control, and Computing*, 2007.

"Improved lower bounds for the capacity of i.i.d. deletion and duplication channels, with E. Drinea. *IEEE Transactions on Information Theory*, vol 53, Issue 8, pp. 2693-2714, 2007.

"Wired Geometric Routing," with J. Ledlie, P. Pietzuch, and M. Seltzer. In *Proceedings of the 6th International Workshop on Peer-to-Peer Systems (IPTPS)*, 2007.

"Capacity Upper Bounds for Deletion Channels," with S. Diggavi and H. Pfister. In *Proceedings of the 2007 Int'l Symposium on Information Theory (ISIT)*, pp. 1716-1720, 2007.

"Codes for Deletion and Insertion Channels with Segmented Errors," with Z. Liu. In *Proceedings of the 2007 Int'l Symposium on Information Theory (ISIT)*, pp. 846-850, 2007. Journal version: *IEEE Transactions on Information Theory*, vol 56, Issue 1, pp. 224-232, 2010.

"Towards a theory of networked computation : Executive Summary," with J. Feigenbaum. *ACM SIGACT News*, (4):22-26, 2006.

"Bloom Filters via d-Left Hashing and Dynamic Bit Reassignment," with F. Bonomi, R. Panigrahy, S. Singh, and G. Varghese. In *Proceedings of the Allerton Conference*, pp. 877-883, 2006.

"Beyond Bloom Filters: From Approximate Membership Checks to Approximate State Machines," with F. Bonomi, R. Panigrahy, S. Singh, and G. Varghese. In *Proceedings of ACM SIGCOMM*, pp. 315-326, 2006.

"An Improved Construction for Counting Bloom Filters," with F. Bonomi, R. Panigrahy, S. Singh, and G. Varghese. In *Proceedings of the European Symposium on Algorithms*, pp. 684-695, 2006.

"Less Hashing, Same Performance: Building A Better Bloom Filter," with A. Kirsch. In *Proceedings of the European Symposium on Algorithms*, pp. 456-467, 2006. Journal version: *Random Structures and Algorithms*, 33(2), pp. 187-218, 2008.

"Stochastic Shortest Paths via Quasi-Convex Maximization," with M. Brand, J. Kelner, and E. Nikolova. In *Proceedings of the European Symposium on Algorithms*, pp. 552-563, 2006.

"Polynomial Time Low-Density Parity-Check Codes with Rates Very Close to the Capacity of the q-ary Random Deletion Channel for Large q." *IEEE Transactions on Information Theory*, vol 52, Issue 12, pp. 5496-5501, 2006.

"A Simple Lower Bound for the Capacity of the Deletion Channel," with E. Drinea. *IEEE Transactions on Information Theory*, vol 52, Issue 10, pp. 4657-4660, 2006.

"On Lower Bounds for the Capacity of Deletion Channels," with E. Drinea. *IEEE Transactions on Information Theory*, vol 52, Issue 10, pp. 4648-4657, 2006.

" On the Theory and Practice of Data Recovery with Multiple Versions." In *Proceedings of the International Symposium on Information Theory*, pp. 982-986, 2006.

"Fine-Grained Layered Multicast with STAIR," with J. Byers, Gu-In Kwon, and M. Luby. *IEEE/ACM Transactions on Networking, 14:1*, pp. 81-93, 2006.

"Distance-Sensitive Bloom Filters," with A. Kirsch. In *Proceedings of the ALENEX Conference*, pp. 41-50, 2006.

"BubbleSearch: A Simple Heuristic for Improving Priority-Based Greedy Algorithms," with N. Lesh. *Information Processing Letters*, 97(4), pp. 161-169, 2006.

"Network-Aware Overlays with Network Coordinates," with P. Pietzuch, J. Ledlie, and M. Seltzer. In *Proceedings of the International Workshop on Dynamic Distributed Systems*, 2006.

"Editorial: The Future of Power Law Research." *Internet Mathematics*, vol. 2. no. 4, pp. 525-534, 2006.

"X-Tolerant Test Response Compaction," with S. Mitra, S. Lumetta, and N. Patil. *IEEE Design and Test of Computers*, vol. 22, no. 6, pp. 566-574, Nov/Dec 2005.

"Privacy preserving keyword searches on remote encrypted data," with Y.-C. Chang. In *Proceedings of the Third Intl. Conference on Applied Cryptography and Network Security*, pp. 442-455, 2005.

"Simple Summaries for Multiple Choice Hashing," with A. Kirsch. In *Proceedings of the 43rd Annual Allerton Conference on Communication, Control, and Computing*, 2005. Journal version: *IEEE/ACM Transactions on Networking*: 16(1), pp. 218-231, 2008.

"A case study in large-scale interactive optimization," with M. Chimani, N. Lesh, and C. Sidner. In *Proceedings of Artificial Intelligence and Applications 2005*, pp. 24-29.

"New heuristic and interactive approaches to 2D rectangular strip packing," with N. Lesh, A. McMahon, and J. Marks. *ACM Journal of Experimental Algorithms*, 10, 2005.

"Multidimensional Balanced Allocations," with A. Broder. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 195-196, 2005.

"Digital Fountains: A Survey and Look Forward." 2004 Information Theory Workshop.

"Improved Lower Bounds on the Capacity of I.I.D. Deletion Channels," with E. Drinea. In *Proceedings of the 42nd Annual Allerton Conference on Communication, Control, and Computing*, 2004.

"Improved Lower Bounds on the Capacity of Deletion Channels," with E. Drinea. In *Proceedings of the 2004 IEEE International Symposium on Information Theory.*

"Geometric Generalizations of the Power of Two Choices," with J. Byers and J. Considine. In *Proceedings of the 16th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 54-63, 2004.

"Interactive Data Summarizaton: An Example Application," with N. Lesh. In *Proceedings of the 2004 Workshop on Advanced Visual Interfaces*, pp. 183-187.

"On Power Law Distributions for Monkeys Typing Randomly: The Case of Unequal Probabilities," with B. Conrad. *IEEE Transactions on Information Theory*, 50:7, pp. 1403-1414, 2004.

"Dynamic Models for File Sizes and Double Pareto Distributions." *Internet Mathematics*, 1(3):305-334, 2003.

"A Scaling Result for Explosive Processes," with R. Oliveira and J. Spencer. *Electronic Journal of Combinatorics*, 11(1), R31, 2004.

"An Exhaustive Approach to 2D Strip Packing," with N. Lesh, A. McMahon, and J. Marks. *Information Processing Letters*, 90(1), pp. 7-14, 2004.

"New Models and Methods for File Size Distributions," with B. Tworetzky. In *Proceedings of the 41st Annual Allerton Conference on Communication, Control, and Computing*, pp. 603-612, 2003.

"Capacity Approaching Signal Constellations for Channels with Memory," with A. Kavčić, and X. Ma. *IEEE Transactions on Information Theory*, 49:7, pp. 1636-1652, 2003.

"The Power Spectra of Good Codes for Partial Response Channels," with A. Kavcic and X. Ma. In *Proceedings of the 2003 IEEE International Symposium on Information Theory*, p. 45.

"Concatenated Codes for Deletion Channels," with J. Chen, C. Ng, and N. Varnica. In *Proceedings of the 2003 IEEE International Symposium on Information Theory*, p. 218.

"Verification Codes for Deletion Channels." In *Proceedings of the 2003 IEEE International Symposium on Information Theory*, p. 217.

"Estimating and Comparing Entropies Across Written Natural Languages Using PPM Compression," with F. Behr, V. Fossum, and D. Xiao. In *Proceedings of the 2003 IEEE Data Compression Conference*, p. 416.

"Simple Load Balancing for Distributed Hash Tables," with J. Byers, and J. Considine. Appears in the 2nd International Workshop on Peer-to-Peer Systems (IPTPS), pp. 80-87, 2003.

"A Complete and Effective Move Set for Simplified Protein Folding," with N. Lesh, and S. Whitesides. In *Proceedings of the 7th Annual International Conference on Research in Computational Molecular Biology*, pp 188-195, 2003.

"Network Applications of Bloom Filters: A Survey," with A. Broder. In *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*, pp. 636-646, 2002. Journal version: *Internet Mathematics*, 1(4):485-509, 2003.

"Verification Codes," with M. Luby. In *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*, pp. 38-47, 2002. Journal version: *IEEE Transactions on Information Theory*: 51(1), pp. 120-127, 2005.

"Balls and Bins with Memory," with B. Prabhakar and D. Shah. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pp. 799-808, 2002.

"Informed Content Delivery Across Adaptive Overlay Networks," with J. Byers, J. Considine, and S. Rost. In *Proceedings of ACM SIGCOMM '02*, available as *Computer Communications Review*, vol. 32:4, pp. 47-60, 2002. Journal version: *IEEE/ACM Transactions on Networking*, 12:5, pp. 767-780, 2004.

"Optimal Plans for Aggregation," with A. Broder. In *Proceedings of the 21st Annual ACM Symposium on Principles of Distributed Computing*, pp. 144-152, 2002.

"Human-Guided Tabu Search," with G. Klau, N. Lesh, and J. Marks. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pp. 41-47, 2002.

"The HuGS Platform: A Toolkit for Interactive Optimization," with G. Klau, N. Lesh, J. Marks, and G. Schafer. In *Proceedings of the Workshop on Advanced Visual Interfaces*, pp. 324-330, 2002.

"Exact Sampling of TCP Window States," with A. Goel. In *Proceedings of IEEE INFOCOM*, pp. 259-265, 2002.

"Balls and Bins Models with Feedback," with E. Drinea and A. Frieze. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 308-315, 2002.

"A Brief History of Lognormal and Power Law Distributions." In *Proceedings of the 39th Annual Allerton Conference on Communication, Control, and Computing*, pp. 182-191, 2001. Journal version: *Internet Mathematics*, 1(2):226-251, 2003.

"Capacity Approaching Signal Constellations for Channels with Memory," with A. Kavčić, X. Ma, and N. Varnica. In *Proceedings of the 39th Annual Allerton Conference on Communication, Control, and Computing*, pp. 311-320, 2001.

"Compressed Bloom Filters." In *Proceedings of the 20th Annual ACM Symposium on Principles of Distributed Computing*, pp. 144-150, 2001. Journal version: *IEEE/ACM Transactions on Networks*, 10:5, pp. 613-620, October 2002.

"Towards More Completed Models of TCP Throughput," with R. Rajaraman. *Journal of Supercomputing*, 20(2), pp. 137-160, September 2001.

"Deriving Performance Bounds for ISI Channels Using Gallager Codes," with A. Kavčić, B. Marcus, and B. Wilson. In *Proceedings of the 2001 International Symposium on Information Theory*, p. 345, 2001.

"Toward Compressing Web Graphs," with M. Adler. In *Proceedings of the 2001 Data Compression Conference*, pp. 203-212, 2001.

"On the Hardness of Finding Multiple Preset Dictionaries." In *Proceedings of the 2001 Data Compression Conference*, pp. 411-418, 2001. Journal version: *IEEE Transactions on Information Theory*, 50:7, pp. 1536-1539, 2004.

"Fine-Grained Layered Multicast," with J. Byers and M. Luby. In *Proceedings of IEEE INFOCOM 2001*, pp. 1143-1151, 2001. Journal version: "Fine-Grained Layered Multicast with STAIR," with J. Byers, Gu-In Kwon, and M. Luby. To appear in *IEEE/ACM Transactions on Networking.*

"Using Multiple Hash Functions to Improve IP Lookups," with A. Broder. In *Proceedings of IEEE INFOCOM 2001*, pp. 1454-1463, 2001.

"Improved Results for Route Planning in Stochastic Transportation Networks," with J. Boyan. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 895-902, 2001.

"Estimating Resemblance of MIDI Documents," with S. Owen. In *Proceedings of the 3rd Workshop on Algorithm Engineering and Experiments*, available as *Lecture Notes in Computer Science 2153*, pp. 78-90, 2001.

"Linear Waste of Best Fit Bin Packing on Skewed Distributions," with Claire Kenyon. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, pp. 582-589, 2000. Journal version: invited to special issue of *Random Structures and Algorithms*, 20(3), pp. 441-464, 2002.

"FLID/DL: Congestion Control for Layered Multicast," with J. Byers, M. Frumin, G. Horn, M. Luby, A. Roetter, and W. Shaver. In *Proceedings of the 2nd International Workshop on Networked Group Communication (NGC 2000)*, pp. 71-81, 2000. Journal version: with J. Byers, G. Horn, M. Luby, and W. Shaver. *IEEE Journal on Selected Areas in Communications*, 20:8, pp. 1558-1570, October 2002.

"On Near-Uniform URL Sampling," with M. Henzinger, A. Heydon, and M. Najork. In *Proceedings of the Ninth International World Wide Web Conference*, pp. 295-308, 2000.

"An Extension of Path Coupling and Its Application to the Glauber Dynamics for Path Coupling," with M. Dyer, L. Goldberg, C. Greenhill, and M. Jerrum. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 616-623, 2000. Journal version: *SIAM Journal on Computing*, 30(6), pp. 1962-1975, 2001.

"Improved Classification via Connectivity Information," with A. Broder and R. Krauthgamer. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 576-585, 2000.

"The Asymptotics of Selecting the Shortest of Two, Improved," with B. Vöcking. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, pp. 326-327, 1999. Extended version appears as a chapter in *Analytic Methods in Applied Probability: In Memory of Fridrikh Karpelevich*, American Mathematical Society, pp. 165-176, 2002.

"Completeness and Robustness Properties of Min-Wise Independent Permutations," with A. Broder. In *Proceedings of Random '99*, available as *Lecture Notes in Computer Science 1671*, pp. 1-10, 1999. Journal version: *Random Structures and Algorithms*, 18(1), pp. 18-30, January 2001.

"Measuring Index Quality Using Random Walks on the Web," with M. Henzinger, A. Heydon, and M. Najork. In *Proceedings of the Eighth International World Wide Web Conference*, pp. 213-225, Elsevier Science, 1999.

"Analysis of Timing Based Mutual Exclusion with Random Times," with E. Gafni. In *Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing*, pp. 13-21, 1999. Journal version: *SIAM Journal on Computing*, 31:3, pp. 816-837, 2001.

"Accessing Multiple Mirror Sites in Parallel: Using Tornado Codes to Speed Up Downloads," with J. Byers and M. Luby. In *Proceedings of the 18th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '99)*, pp. 275-284, 1999.

"Unscrambling Address Lines (Abstract)," with A. Broder and L. Moll. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 870-871, 1999.

"A Derandomization Using Min-Wise Independent Permutations," with A. Broder and M. Charikar. In *Proceedings of Random '98*, available as *Lecture Notes in Computer Science 1518*, pp. 15-24, October 1998. Journal version: *Journal of Discrete Algorithms*, 1:1, pp. 11-20, 2003.

"On Balls and Bins with Deletions," with R. Cole, A. Frieze, B. M. Maggs, A. W. Richa, R. Sitaraman, and E. Upfal. In *Proceedings of Random '98*, available as *Lecture Notes in Computer Science 1518*, pp. 145-158, October 1998.

"Studying Balanced Allocations with Differential Equations." *Combinatorics, Probability, and Computing*, vol. 8, pp. 473-482, 1999.

"Delayed Information and Action in On-Line Algorithms," with S. Albers and M. Charikar. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science (FOCS-98)*, pp. 71-80, November 1998. Journal version: *Information and Computation*, 170, pp. 135-152, 2001.

"A Digital Fountain Approach to the Reliable Distribution of Bulk Data," with John Byers, Michael Luby, and Ashu Rege. In *Proceedings of ACM SIGCOMM '98*, available as *Computer Communications Review*, vol. 28:4, pp. 56-67, 1998. Journal version: *IEEE Journal on Selected Areas in Communication*, 20:8, pp. 1528-1540, October 2002.

"Analyses of Load Stealing Models Using Differential Equations." In *Proceedings of the 10th Annual ACM Symposium on Parallel Algorithms and Architectures*, pp. 212-221, 1998. Journal version: *Theory of Computing Systems*, 34:1, pp. 77-98, January 2001.

"Load Adaptive Routing on Butterfly Networks," with R. Cole, B. Maggs, F. Meyer auf der Heide, A. Richa, K. Schröder, R. Sitaraman, and B. Vöcking. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)*, pp. 378-388, ACM Press, May 23-26, 1998.

"Min-Wise Independent Permutations," with Andrei Broder, Moses Charikar, and Alan Frieze. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)*, pp. 327-336, ACM Press, May 23-26, 1998. Journal version: *Journal of Computer and System Sciences*, 60(3), pp. 630-659, 2000.

"Analysis of Low Density Codes and Improved Designs Using Irregular Graphs," with Michael Luby, Amin Shokrollahi, and Dan Spielman. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)*, pp. 249-258, ACM Press, May 23-26, 1998.

"Improved Low Density Parity Check Codes Using Irregular Graphs and Belief Propagation," with Michael Luby, Amin Shokrollahi, and Dan Spielman. In *Proceedings of the 1998 International Symposium on Information Theory.* Extended version available as Compaq Systems Research Center Technical Note 1998-009. Journal version: "Improved Low Density Parity Check Codes Using Irregular Graphs," *IEEE Transactions on Information Theory*, 47(2), pp. 585-598, 2001.

"Analysis of Random Processes via And-Or Tree Evaluation," with Michael Luby and Amin Shokrollahi. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 364-373, 1998.

"Average Case Analyses of First Fit and Random Fit Bin Packing," with Susanne Albers. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp.

290-299, 1998. Journal version: *Random Structures and Algorithms*, 16(3), pp. 240-259, 2000.

"How Useful is Old Information?" In *Proceedings of the 16th ACM Symposium on Principles of Distributed Computing*, pp. 83-91, 1997. Journal version: *IEEE Transactions on Parallel and Distributed Systems*, 11:1, pp. 6-20, January 2000.

"Revisiting the COUNTER Algorithms for List Update," with Susanne Albers. *Inf. Proc. Letters*, 64, pp. 155-160, 1997.

"Practical Loss-Resilient Codes," with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. In *Proceedings of the 29th ACM Symposium on the Theory of Computing*, pp. 150-159, 1997. Journal version: "Efficient Erasure Correcting Codes," with Michael Luby, Amin Shokrollahi, and Dan Spielman. *IEEE Transactions on Information Theory*, 47(2), pp 569-584, 2001.

"On the Analysis of Randomized Load Balancing Schemes." In *Proceedings of the 9th ACM Symposium on Parallel Algorithms and Architectures*, pp. 292-301, 1997. Journal version: special conference issue of *Theory of Computing Systems*, vol. 32, pp. 361-386, 1999.

"Load Balancing and Density Dependent Jump Markov Processes." In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pp. 213-222, 1996. Journal version: "The Power of Two Choices in Randomized Load Balancing." *IEEE Transactions on Parallel and Distributed Computing*, 12(10), pp. 1094-1104, 2001.

"Designing Stimulating Programming Assignments for an Algorithms Course: A Collection of Problems Based on Random Graphs." In *SIGCSE Bulletin*, 28:3, pp. 29-36, September 1996.

"Constant Time per Edge is Optimal in Rooted Tree Networks." In *Proceedings of the 8th ACM Symposium on Parallel Algorithms and Architectures*, pp. 162-169, 1996. Journal version: *Distributed Computing* vol. 10:4, 1997, pp. 189-197.

"Average Case Analyses of List Update Algorithms, with Applications to Data Compression," with Susanne Albers. In *Proceedings of the 1996 International Colloquium on Automata, Languages, and Programming*, pp. 514-525. Journal version: *Algorithmica*, vol. 21:3, 1998, pp. 312-329.

"Pattern-based Compression of Text Images," with Andrei Broder. In *1996 Data Compression Conference*, pp. 300-309.

"Parallel Randomized Load Balancing," with Micah Adler, Soumen Chakrabarti, and Lars Rasmussen. In *Proceedings of the 27th ACM Symposium on the Theory of Computing*, pp. 238-247, 1995. Journal version: *Random Structures and Algorithms*, vol. 13:2, September 1998, pp. 159-188.

"Bounds on the Greedy Routing Algorithm in Array Networks." In *Proceedings of the 6th ACM Symposium on Parallel Algorithms and Architectures*, pp. 346-353, 1994. Journal version: special conference issue of the *Journal of Computer and System Sciences* vol. 53:3, December 1996, pp. 317-327.

"Computational Complexity of Loss Networks," with Graham Louth and Frank Kelly. *Theoretical Computer Science*, vol. 125, May 1994, pp. 45-59.

Other Publications

"Lyra: An Efficient and Expressive Subquadratic Architecture for Modeling Biological Sequences," with Krithik Ramesh, Sameed Siddiqui, Albert Gu, and Pardis Sabeti. Available at `https://arxiv.org/abs/2503.16351`.

"Juniper: Reconstructing Transmission Events from Next-Generation Sequencing Data at Scale," with Ivan Specht, Gage Moreno, Taylor Brock-Fisher, Lydia Kraslinkova, Brittany Petros, Jonathan Pekar, Mark Schifferli, Ben Fry, Catherine Brown, Lawrence Madoff, Megan Burns, Stephen Schaffner, Daniel Park, Bronwn MacInnis, Al Ozonoff, Patrick Varilly, and Pardis Sabeti. Available at `https://www.medrxiv.org/content/10.1101/2025.03.02.25323192v1`.

"Delphy: scalable, near-real-time Bayesian phylogenetics for outbreaks," with Patrik Varilly, Mark Schifferli, Katherine Yang, Tim Burcham, Paul Cronan, Olivia Glennon, Olivia Jacks, Ellory Laning, Libby Marrs, Kyle Oba, Shannon Yeung, Edyth Parker, Ifeanyi Omah, Jonathan E. Pekar, Laura Luebbert, Kristian G. Andersen, Daniel J. Park, Stephen F. Schaffner, Bronwyn L. MacInnis, Christian Happi, Jacob E. Lemieux, Al Ozonoff, Michael D. Mitzenmacher, Ben Fry, and Pardis C. Sabeti. Available at `https://www.biorxiv.org/content/10.1101/2025.03.25.645253v1`.

"Use of popularity information to reduce risk posed by guessing attacks," with Stuart Schechter and Cormac Herley. US Patent 08959644. Issued 2/17/2015.

"Method and apparatus providing network coding based flow control," with Jay Kumar Sundararajan, Devavrat Shah, Muriel Medard, Joao Barros, and Szymon Jakubczak. US Patent 08526451. Issued 9/3/2013.

"Method and apparatus providing network coding based flow control," with Jay Kumar Sundararajan, Devavrat Shah, Muriel Medard, Joao Barros, and Szymon Jakubczak. US Patent 08130776. Issued 3/6/2012.

"Method for finding optimal paths using a stochastic network model," with Evdokia Nikolova and Matthew Brand. US Patent 07573866. Issued 8/11/2009.

"Ranking search engine results," with Monika Henzinger. US Patent 07541388. Issued 11/11/2008.

"Distributed, compressed Bloom filter Web cache server." US Patent 06920477. Issued 7/19/2005.

"Method for packing rectangular strips," with Neal Lesh and Joe Marks. US Patent 06832129. Issued 12/14/2004.

"System and method for near-uniform sampling of web page addresses," with Monika Henzinger, Allan Heydon, and Marc Najork. US Patent 06594694. Issued 7/15/2003.

"On demand encoding with a window," with Michael Luby, Gavin Horn, Jeffrey Persch, John Byers, and Armin Haken. US Patent 06486803. Issued 11/26/2002.

"Generating high weight encoding symbols using a basis," with Armin Haken, Michael Luby, Gavin Horn, Diane Hernek, and John Byers. US Patent 06411223. Issued 6/25/2002.

"Space-efficient multi-cycle barrel shifter circuit," with Laurent Moll. US Patent 06314156. Issued 11/6/2001.

"Method for determining a random permutation of variables by applying a test function," with Andrei Broder, Laurent Moll, and Mark Shand. US Patent 06292762. Issued 9/18/2001.

"Loss resilient code with double heavy tailed series of redundant layers," with Michael Luby. US Patent 06195777. Issued 2/27/2001.

"Message encoding with irregular graphing," with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. US Patent 06163870. Issued 12/19/2000.

"Compression of grey scale images of text," with Andrei Broder. US Patent 06088309. Issued 7/11/2000.

"Irregularly graphed encoding technique," with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. US Patent 06081909. Issued 6/27/2000.

"Loss resilient decoding technique," with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. US Patent 06073250. Issued 6/6/2000.

"Task processing optimization in a multiprocessor system," with Andrei Broder. US Patent 05991808. Issued 11/23/1999.

"Compression protocol with multiple preset dictionaries," with Andrei Broder and Jeff Mogul. US Patent 05953503. Issued 9/14/1999.

"When Simple Hash Functions Suffice," with K. Chung and S. Vadhan. Appears as a chapter in *Beyond Worst Case Analysis*, edited by Tim Roughgarden, Cambridge University Press.

"Algorithms with Predictions," with S. Vassilvitskii. Appears as a chpater in *Beyond Worst Case Analysis*, edited by Tim Roughgarden, Cambridge University Press.

"Some Practical Randomized Algorithms and Data Structures," with E. Upfal. Appears as a chapter in *Computing Handbook, Third Edition: Computer Science and Software Engineering*, edited by T. Gonzalez, J. Diaz-Herrera, and A. Tucker. Chapman and Hall, 2014.

"Building an Efficient Hash Table on the GPU," with D. Alcantara, V. Volkov, S. Sengupta, J. Owens, and N. Amenta. Appears as a chapter in *GPU Computing Gems: Jade Edition*, edited by Wen-mei W. Hwu. Morgan Kaufmann, 2011.

"Codes – Protecting Data Against Loss and Errors." Appears as a chapter in *Algorithms Unplugged*, edited by B. Vöcking, M. Dietzfelbinger, R. Reischuk, and H. Vollmer. Springer, 2010.

"Hash-Based Techniques for High-Speed Packet Processing", with A. Kirsch and G. Varghese. Appears as a chapter in *Algorithms for Next Generation Networks*, edited by G. Cormode and M. Thottan. Springer, 2010.

"Bloom Filters." Encyclopedia of Database Systems, 2009.

"Unbiasing Random Bits." *Dr. Dobbs Journal*, No. 335, pp. 101-104, April 2002.

"Good Hash Tables & Multiple Hash Functions." *Dr. Dobbs Journal*, No. 336, May 2002.

"Challenging Students with Creative Assignments." *SIGACT News*, 32(1):70-73, 2001.

"An Experimental Assignment on Random Processes." *SIGACT News*, 32(1):74-78, 2001.

"The Power of Two Random Choices: A Survey of Techniques and Results," with A. Richa and R. Sitaraman. Appears as a chapter in *Handbook of Randomized Computing: Volume 1*, edited by P. Pardalos, S. Rajasekaran, and J. Rolim, pp. 255-312.

"Tornado Codes for Forward Error Correction," with Michael Luby, in *Digital Forefront.*

"Bounds and Improvements for BiBa Signature Schemes," with A. Perrig. Harvard Computer Science Group Technical Report TR-02-02.

"A Note on Low Density Parity Check Codes for Erasures and Errors." Digital SRC Technical Note 1998-017.

"Tight Thresholds for the Pure Literal Rule." Digital SRC Technical Note 1997-011.

"The Power of Two Choices in Randomized Load Balancing." Ph.D. thesis, U.C. Berkeley, August 1996.

"*Calculus: A Solution Manual,*" with others, Harcourt Brace Jovanovich, 1993.

"*Elliptic Curves in Computer Science: Primality Testing, Factoring, and Cryptography,*" senior honors thesis, Harvard University, 1991. Published as a Hoopes Prize winner.

Conference Presentations

"Queueing, Predictions, and LLMs: Challenges and Open Problems." Tutorial, SIGMETRICS. June 2025.

"SkipPredict: When to Invest in Predictions for Scheduling. Presented at the TTIC Workshop on Learning-Augmented Algorithms, August 2024.

"Algorithms with Predictions." Invited talk presented at the INFORMS 2023 Annual Meeting, October 2023.

"Algorithms with Predictions." Presented at the 2nd Vienna Workshop on Computational Optimization. March 2023.

"Algorithms with Predictions." Presented at the Sublinear Algorithms Workshop, FODSI/MIT. August 2022.

"Algorithms with Predictions: State of the Area." Presented at the Algorithms with Predictions Workshop at the 54th Annual ACM Symposium on Theory of Computing. June 2022. Also panel lead.

"Efficient Randomized Network Telemetry Algorithms." Presented at the Data plane algorithms in programmable networks tutorial at ACM SIGMETRICS 2022.

"Scheduling with Predictions and the Price of Misprediction." Presented at the 11th Annual Innovations in Theoretical Computer Science Conference. January 2020.

"Learning-Augmented Algorithms: How ML Can Lead to Provably Better Algorithms." ALGO Keynote Talk, September 2019.

"Robust Set Reconciliation via Locality Sensitive Hashing." Presented at the 38th Symposium on Principles of Database Systems, July 2019.

"Arithmetic Progression Hypergraphs: Examining the Second Moment Method." Presented at the 2019 ANALCO Workshop, January 2019.

"A Model for Learned Bloom Filters and Optimizing by Sandwiching." Presented at the NeurIPS 2018 Poster Session.

"Adaptive Cuckoo Filters." Presented at the 2018 ALENEX Conference, January 2018.

"Analyzing Distributed Join-Idle-Queue: A Fluid Limit Approach." Presented at the 54th Annual Allerton Conference on Communication, Control, and Computing, October 2016.

"Voronoi Choice Games." Presented at the 2016 International Colloquium on Automata, Languages, and Programming, July 2016.

"A New Approach to Analyzing Robin Hood Hashing." Presented at the 2016 ANALCO Workshop, January 2016.

"More Analysis of Double Hashing for Balanced Allocations." Presented at the 2016 ANALCO Workshop, January 2016.

"Data exploration: equitability and the maximal information coefficient." Presented at the Workshop for Data-driven Algorithmics, September 2015.

"Multi-Party Set Reconciliation Using Characteristic Polynomials." Presented at the 2014 Allerton Conference, October 2014.

"Parallel Peeling Algorithms." Presented at the 2014 Symposium on Parallelism in Algorithms and Architectures, June 2014.

"Balanced Allocations and Double Hashing." Presented at the 2014 Symposium on Parallelism in Algorithms and Architectures, June 2014.

"Peeling Arguments and Double Hashing." Presented at the 2012 Allerton Conference, October 2012.

"The Complexity of Object Reconciliation and Open Problems Related to Set Difference and Coding." Presented at the 2012 Allerton conference, October 2012.

"The Ecology of Communication, Control, and Computing Research and Education (Panel)." Presented at the 2012 Allerton conference, October 2012.

"Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes." Presented at the 2012 Workshop on Algorithms for Modern Massive Data Sets (MMDS), July 2012.

"Invertible Bloom Lookup Tables." Presented at the 2011 Allerton conference, September 2011.

"Brief Announcement: Large-Scale Multimaps." Presented at the 2011 Symposium on Parallel Algorithms and Architectures, June 2011.

"On the Zero-Error Capacity Threshold for Deletion Channels." Presented at the 2011 ITA Workshop, February 2011.

"Some Recent Uses of Hashing for Network Applications." Presented at the 2010 Workshop on Research and Analysis of Tail Phenomena, August 2010.

"Some Recent Uses of Hashing for Network Applications." Presented at the 2010 Workshop on Random Graphs and Algorithms, June 2010.

"Tight Asymptotic Bounds for the Deletion Channel with Small Deletion Probabilities." Presented at the 2010 International Symposium on Information Theory, June 2010.

"Some Open Questions on Cuckoo Hashing." Presented at the 2009 European Symposium on Algorithms, August 2009.

"On the Performance of Multiple Choice Hash Tables with Moves on Deletions and Inserts." Presented at the 2008 Allerton conference, September 2008.

"Designing Floating Codes for Expected Performance." Presented at the 2008 Allerton conference, September 2008.

"A Survey of Results for Deletion Channels and Related Synchronization Channels.' Presented at the 11th Scandinavian Workshop on Algorithm Theory, July 2008.

"Why Simple Hash Functions Work: Exploiting the Entropy in a Data Stream." Presented at the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2008.

"Using a Queue to De-amortize Cuckoo Hashing in Hardware." Presented at the 2007 Allerton conference, September 2007.

"Capacity Upper Bounds for Deletion Channels." Presented at the 2007 Int'l Symposium on Information Theory (ISIT), June 2007.

"Codes for Deletion and Insertion Channels with Segmented Errors." Presented at the 2007 Int'l Symposium on Information Theory (ISIT), June 2007.

"Digital Fountains: Applications and Related Issues." Presented at the INFORMS Annual Meeting, November 2006.

"New Directions for Power Law Research." Presented at the Radcliffe Exploratory Seminar on Dynamic Networks, October 2006.

"Bloom Filters via d-left Hashing and Dynamic Bit Reassignment." Presented at the 2006 Allerton conference, September 2006.

"Beyond Bloom Filters: From Approximate Membership Checks to Approximate State Machines." Presented at the 2006 SIGCOMM conference, September 2006.

"An Improved Construction for Counting Bloom Filters." Presented at the European Symposium on Algorithms, September 2006.

"Building a Better Bloom Filter." Presented at the European Symposium on Algorithms, September 2006.

"On the Theory and Practice of Data Recovery with Multiple Versions." Presented at the 2006 International Symposium on Information Theory, July 2006.

"New Results for an Old Data Structure: Bloom Filters and Beyond." Presented at the 2006 Stochastic Networks conference, June 2006.

"New Results and Open Problems for Deletion Channels." Presented at ANALCO 2006, January 2006 (plenary talk).

"Digital Fountains and their Application to Informed Content Delivery over Adaptive Overlay Networks." Presented at 19th International Symposium on Distributed Computing, September 2005 (plenary talk).

"Simple Summaries for Multiple Choice Hashing." Presented at the 2005 Allerton Conference.

"New Directions in Power Law Research." MSRI Workshop on Models of Real-World Random Networks, April 2005.

"Multi-dimensional Balanced Allocations." Presented at the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2005.

"Digital Fountains: A Survey and Look Forward." Presented at the 2004 Information Theory Workshop. October 2004.

"Improved Lower Bounds on the Capacity of I.I.D. Deletion Channels." Presented at the 2004 Allerton Conference.

"Geometric Generalizations of the Power of Two Choices." Presented at the 16th ACM Symposium on Parallel Algorithms and Architectures, June 2004.

"New Models and Methods for File Size Distributions." Presented at the 40th Annual Allerton Conference on Communication, Control, and Computing, October 2003.

"Verification Codes for Deletion Channels." Presented at the International Symposium on Information Theory, July 2003.

"New Exhaustive, Heuristic, and Interactive Algorithms for 2-Dimensional Bin-Packing." Presented at the 2nd Annual Workshop on Random Graphs and Algorithms, June 2003.

"Network Applications of Bloom Filters: A Survey." Presented at the 40th Annual Allerton Conference on Communication, Control, and Computing, October 2002.

"Verification Codes." Presented at the 40th Annual Allerton Conference on Communication, Control, and Computing, October 2002.

"New Directions in Balls and Bins Problems." Invited talk for the 6th International Workshop on Randomization and Approximation Methods in Computer Science. September 2002.

"Optimal Plans for Aggregation." Presented at the 21st Annual ACM Symposium on Principles of Distributed Computing, July 2002.

"How to Get Across Town by Bus." Presented at the INFORMS Annual Meeting, November 2001.

"A Brief History of Lognormal and Power Law Distributions." Presented at the 39th Annual Allerton Conference on Communication, Control, and Computing, October 2001.

"Compressed Bloom Filters." Presented at the 20th Annual ACM Symposium on Principles of Distributed Computing, August 2001.

"How Useful Is Old Information?" Presented at the INFORMS Applied Probability Conference, July 2001.

"Using Multiple Hash Functions to Improve IP Lookups." Presented at the INFOCOM, April 2001.

"On the Hardness of Finding Multiple Preset Dictionaries." Presented at the 2001 Data Compression Conference, March 2001.

"Improved Results for Route Planning in Stochastic Transportation Networks." Presented at the 12th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2001.

"Estimating Resemblance of MIDI Documents." Presented at the 3rd Workshop on Algorithm Engineering and Experiments, January 2001.

"Linear Waste of Best Fit Bin Packing on Skewed Distributions." Presented at the 41st Annual IEEE Symposium on Foundations of Computer Science, November 2000.

"Improved Classification via Connectivity Information." Presented at the 11th Annual ACM Symposium on Discrete Algorithms, January 2000.

"The Asymptotics of Selecting the Shortest of Two, Improved." Presented at the 37th Annual Allerton Conference on Communication, Control, and Computing, August 1999.

"Analysis of Timing Based Mutual Exclusion with Random Times." Presented at the 18th Annual ACM Symposium on Principles of Distributed Computing, May 1999.

"Accessing Multiple Mirror Sites in Parallel: Using Tornado Codes to Speed Up Downloads." Presented at INFOCOM '99, March 1999.

"Delayed Information and Action in On-Line Algorithms." Presented at the 39th Annual IEEE Symposium on Foundations of Computer Science, November 1998.

"A Derandomization Using Min-Wise Independent Permutations." Presented at Random '98, October 1998.

"On Balls and Bins with Deletions." Presented at Random '98, October 1998.

"The Power of Two Choices and Other Examples of Using Differential Equations to Analyze Algorithms." Invited presentation at the Warwick Randomized Algorithms and Stochastic Simulation Tutorial and Workshop. July 1988.

"Analysis of Low Density Codes and Improved Designs Using Irregular Graphs." Presented at the 30th Annual ACM Symposium on Theory of Computing, May 1998.

"Average Case Analyses of First Fit and Random Fit Bin Packing." Presented at the 9th Annual ACM Symposium on Discrete Algorithms, January 1998.

"How Useful Is Old Information?" Presented at the 16th Annual ACM Symposium on Principle of Distributed Computing, August 1997.

"On the Analysis of Randomized Load Balancing Schemes." Presented at the 9th ACM Symposium on Parallel Algorithms and Architectures, Newport, Rhode Island, June 1997.

"Load Balancing and Density Dependent Jump Markov Processes." Presented at the 37th Annual Symposium on Foundations of Computer Science, Burlington, Vermont, October 1996.

"Constant Time per Edge is Optimal in Rooted Tree Networks." Presented at the 8th ACM Symposium on Parallel Algorithms and Architectures, Padua, Italy, June 1996.

"GRID: Pattern-based Compression of Text Images." Presented at the 1996 Data Compression Conference, Snowbird, Utah, April 1996.

"Parallel Randomized Load Balancing." Presented at the 27th ACM Symposium on the Theory of Computing, Las Vegas, Nevada, May 1995. Extended version presented at the Bay Area Theory Seminar, July 1995.

"Bounds on the Greedy Routing Algorithm on Array Networks." Presented at the 6th ACM Symposium on Parallel Algorithms and Architectures, Cape May, New Jersey, June 1994.

Other Presentations

"Algorithms with Predictions in Scheduling." Johns Hopkins University. 2025.

"Algorithms with Predictions." Georgia Tech. 2023.

"Modern Algorithms Lecture Series." Lipari Summer School. 2022.

"Algorithms with Predictions." Carnegie Mellon, UC Santa Barbara, Technical University of Munich. 2022.

"Locality Sensitive Hashing." Broad Institute (Models, Inference, and Algorithms).

"Algorithms with Predictions." Open Data Science Conference East 2020, Open Data Science Conference Europe 2020, Foundations of Data Science Institute Kickoff Workshop 2020.

"Learning-Augmented Algorithms: How ML Can Lead to Provably Better Algorithms." Rice University. 2019.

"Metric Sublinear Algorithms via Linear Sampling." Sublinear Algorithms Workshop (MIT), Baidu Research. 2018.

"Bloom Filters, Cuckoo Hashing, Cuckoo Filters, and Adaptive Cuckoo Filters." Northeastern University, Univ. of Illinois at Urbana-Champagne, Georgetown Univeristy, University of California Irvine, Baidu Research, Georgia Tech. 2018.

"Hashing Lecture Series." Swedish Summer School in Computer Science. 2016.

"More Analysis of Double Hashing for Balanced Allocations." U.C. Irvine. 2016.

"Algorithm Engineering with Hashing," U.C. Irvine, AMD. 2016.

"Bloom Filters and Such," EADS Summer School on Hashing: Theory and Applications, July 2014.

"Cuckoo Hashing, Balanced Allocations, and Such," EADS Summer School on Hashing: Theory and Applications, July 2014.

"Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes," ICERM Workshop on Stochastic Graph Models, March 2014.

"Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes," University of Copenhagen. 2013.

"Some Hash-Based Data Structures and. Algorithms Everyone Should Know." UC Davis, 2013.

"Peeling Algorithms," Simons Institute. 2013.

"Verifying Computations in the Cloud (and Elsewhere)," TTI/Vanguard : Ginormous Systems. Microsoft Faculty Summit. 2013.

"The Groupon Effect on Yelp Ratings: a Root Cause Analysis," Cornell, 2012.

"Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes," Google. 2012.

"Detecting Novel Associations in Large Data Sets." eHarmony (LA Machine Learning). 2012.

"Cuckoo Hashing: New Results and Open Questions." University of Madison-Wisconsin. 2011.

'Practical Verified Computation with Streaming Interactive Proofs." University of Madison-Wisconsin. 2011.

"The Power of Choice: Simple Hash-Based Data Structures Everyone Should Know." Santa Fe Institute. 2011.

"Information Asymmetries in Pay-Per-Bid Auctions: How Swoopo Makes Bank." Liverpool University, Microsoft UK. 2010.

"Some Open Questions on Cuckoo Hashing." Stanford, Microsoft Research, Yale, Dartmouth, Princeton, Cambridge University, University College London . 2009-2010.

"An Efficient, Rigorous Approach for Identifying Statistically Significant Frequent Itemsets." Google. 2009.

"Some Recent Results on Cuckoo Hashing." AT&T. 2009.

"Some Results on Coding for Flash Memory." University College London. 2009.

"Network Coding Meets TCP." University College London. 2009.

"A Survey of Results for Deletion Channels and Related Synchronization Channels." Cambridge, Georgia Tech, Microsoft Research, UCLA, 2009-2010.

"Bloom Filters, Related Data Structures, and their Applications." Cambridge (Kuwait Lecture). 2009.

"A History of and New Directions for Power Law Research." SUNY Buffalo. 2009.

"Why Simple Hash Functions Work." Cisco. 2008.

"Cuckoo Hashing and CAMs." Cisco, Google. 2008.

"New Results in Hash-Based Data Structures." Cisco. 2007.

"The Hiring Problem and Lake Wobegon Strategies." Microsoft, U.C. Berkeley, Princeton, Google, Yahoo. 2007-2008.

"A Brief History of Lognormal and Power Law Distributions." Harvard Statistics. 2006.

"New Results and Open Problems for Deletion Channels and Related Synchronization Problems." CMU, Brown, Harvard. 2006.

"Some New Results on an Old Data Structure, Bloom Filters ." Stanford, Cisco, Microsoft. 2005.

"Improved Lower Bounds for Deletion Channels." MIT, EPFL (Lausanne). 2005.

"Building a Better Bloom Filter." Google, Tufts, EPFL (Lausanne), MIT. 2004-05.

"Geometric Generalizations of the Power of Two Choices." Columbia, Brown, MIT, Georgia Tech. 2004-05.

"New Lower Bounds on the Capacity of the Binary Deletion Channel." MIT, 2004.

"Digital Fountains, and their Use for Informed Content Delivery across Overlay Networks." Cornell, 2003.

"Verification Codes for Error and Deletion Channels." U.C. Berkeley, 2003.

"New Exhaustive, Heuristic, and Interactive Algorithms for 2D Strip Packing." U.C. Berkeley, 2003.

"Dynamic Models for File Sizes and Double Pareto Distributions." IBM Research, Microsoft Silicon Valley, Stanford, Tufts, Boston University, Random Graphs and Algorithms Workshop. 2002.

"Verification Codes." DIMACS workshop. 2001.

"Getting Across Town by Bus, Aggregation, and Reliability." MIT (LIDS), Mitsubishi Electronic Research Labs. 2001.

"New Ideas for Layered Multicast Congestion Control." International Computer Science Institute. 2001.

"Compressed Bloom Filters, and Towards Compressing Web Graphs." U. of Washington, AT&T Labs, U. C. Berkeley, MIT, Brandeis, Carnegie Mellon University, Stanford, DIMACS workshop on compression, U.C. San Diego. 2001.

"Toward Compressing Web Graphs." Compaq Systems Research Center, Altavista Corporation. 2000.

"Improved Results for Route Planning in Stochastic Transportation Networks." Compaq Systems Research Center. 2000.

"Choosing the Shortest of Two, Improved, with Applications to IP Routing." U. Washington, Princeton, Yale, U.C. Berkeley, Boston University, Stanford, M.I.T., Northeastern, Compaq Systems Research Center. 2000.

"Tornado Codes, with Applications to Reliable Multicast." Harvard, CMU, Duke, AT&T Labs, Univ. of Illinois at Urbana-Champagne, University of Rhode Island, Boston University. 1998-2000.

"How Useful Is Old Information?" MIT, U.C. Berkeley, Stanford, Microsoft. 1997.

"Practical Loss-Resilient Codes." U.C. Berkeley, Harvard. 1997.

"An Average Case Analysis for First-Fit Bin Packing." Stanford, DIMACS workshop. 1997.

"The Power of Two Choices." U.C. Berkeley, Brown, Harvard, Boston University, IBM Almaden, Digital Systems Research Center. 1997.

"The Power of Two Choices and Other Examples of Using Differential Equations to Analyze Algorithms." U. Mass.