

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE PATENT TRIAL AND APPEAL BOARD

---

ACTIVISION BLIZZARD, INC.,  
ELECTRONIC ARTS INC.,  
TAKE-TWO INTERACTIVE SOFTWARE, INC.,  
2K SPORTS, INC., and  
ROCKSTAR GAMES, INC.,  
Petitioners,

v.

ACCELERATION BAY, LLC,  
Patent Owner.

---

Case No. IPR2016-00747  
Patent 6,732,147

---

**DECLARATION OF MICHAEL GOODRICH, Ph.D. IN SUPPORT OF  
PATENT OWNER'S RESPONSE TO PETITION**

I, Michael Goodrich, Ph.D., declare and state as follows:

**I. QUALIFICATIONS**

1. I make this Declaration based upon my own personal knowledge, information, and belief, and I would and could competently testify to the matters set forth herein if called upon to do so.

2. I received a Bachelor of Arts (“BA”) degree in Mathematics and Computer Science from Calvin College in 1983 and a PhD in Computer Sciences from Purdue University in 1987.

3. I am a Chancellor's Professor in the Department of Computer Science at the University of California, Irvine, where I have been a faculty member since 2001. The Chancellor's Professor title at University of California, Irvine is designed for persons who have earned the title of Professor and who have demonstrated unusual academic merit and whose continued promise for scholarly achievement is unusually high. In addition, I am technical director for the Center for Algorithms and Theory of Computation in the Donald Bren School of Information and Computer Sciences at University of California, Irvine. I was a professor in the Department of Computer Science at Johns Hopkins University from 1987-2001.

4. I have authored or coauthored over 300 publications, including several widely adopted books, such as *Introduction to Computer Security* and *Algorithm*

*Design and Applications.* My research includes contributions to data structures and algorithms, information security and privacy, networking, graph algorithms, computational geometry, distributed and parallel algorithms, and cloud security. My research is supported by grants from the Defense Advanced Research Projects Agency (DARPA) and the National Science Foundation (NSF).

5. In addition, I have consulting experience in matters involving algorithms, cryptography, machine learning, digital rights management, computer security, networking, software, and storage technologies.

6. I am an ACM Distinguished Scientist, a Fellow of the American Association for the Advancement of Science (AAAS), a Fulbright Scholar, a Fellow of the Institute of Electrical and Electronics Engineers (IEEE), and a Fellow of the Association for Computing Machinery (ACM). I am also a recipient of the IEEE Computer Society Technical Achievement Award and the Pond Award for Excellence in Undergraduate Teaching.

7. Attached hereto as Appendix A is a true and correct copy of my Curriculum Vitae (CV).

8. In developing my opinions below, I have applied my professional experience and I have considered the materials cited herein, including the subject Petition, all the exhibits cited therein and identified on Petitioner's Exhibit List, and the Decision to Institute Trial. In addition, I have reviewed the documents

cited by Dr. David R. Karger in his declaration to the Petition (Dr. Karger's declaration hereinafter referred to as "Karger Decl"). I have also review Dr. Karger's deposition transcript in this case. (Dr. Karger's transcript hereinafter referred to as "Karger Tr." (Ex. 2010)).

9. To the extent that I am presented with new information concerning the subject matter of this declaration or affecting any assumptions made herein, I reserve the right to supplement this declaration accordingly.

## **II. SCOPE OF ASSIGNMENT AND APPROACH**

10. I have been retained as an expert on behalf of Patent Owner, Acceleration Bay LLC, ("Acceleration Bay"), to provide information and opinions to the Patent Trial and Appeal Board (hereinafter "the Board") to assist in the determination of whether or not the Board should invalidate of one or more claims of Acceleration Bay's U.S. Patent No. 6,732,147 (Ex. 1001, the "'147 Patent"). Specifically, counsel for Acceleration Bay asked me to provide opinions regarding the validity of claims 6-10 of the '147 Patent in view of certain prior art references cited by Activision Blizzard, Inc., Electronic Arts Inc., Take-Two Interactive Software, Inc., 2K Sports, Inc., and Rockstar Games, Inc. (collectively, "Petitioners"). In my opinion, claims 6-10 of the '147 Patent are valid in light of the prior art and arguments presented in the Petition and the Karger Decl.

11. I have been informed by counsel and I understand that the analysis of whether a patent is anticipated or obvious is performed from the perspective of a person of ordinary skill in the art at the time of the patented inventions. The relevant timeframe for the challenged claims of the '147 Patent is at least July 2000.

### **III. APPLICABLE STANDARDS AND CONTROLLING PRINCIPLES**

#### **A. Anticipation**

12. Counsel has informed me, and I understand, that an issued patent claim is invalid as anticipated if each and every element of that claim is disclosed in a single prior art reference that enables a person of ordinary skill in the art to make the allegedly anticipating subject matter. I understand that to be anticipatory, a reference must enable one of skill in the art to practice an embodiment of the claimed invention without undue experimentation.

13. Counsel has informed me, and I understand, that if a prior art reference does not disclose a given element expressly, it may do so inherently. I have been informed by counsel and I further understand that a prior art reference will inherently anticipate a claimed invention if any claim elements or other information missing from the reference would nonetheless be known by the person of ordinary skill in the art to be necessarily present in the subject matter of the reference.

**B. Obviousness**

14. Counsel has informed me, and I understand, that an issued patent claim is invalid as obvious if it can be shown that the differences between the patented subject matter and the prior art are such that the subject matter as a whole would have been obvious, at the time the invention was made, to a person having ordinary skill in the art. Relevant considerations include the level of ordinary skill in the art; the scope and content of the prior art; differences between the prior art and the claims at issue; and the so-called objective secondary factors of nonobviousness.

15. Counsel has informed me, and I understand, that in order to evaluate the obviousness of any claim of the '147 Patent over a given prior art combination, I should analyze whether the prior art references, included collectively in the combination, disclose each and every element of the allegedly invalid claim as those references are read by the person of ordinary skill in the art at the time of the invention. Then I am to determine whether that combination makes the claims of the '147 Patent obvious to the person of ordinary skill in the art by a preponderance of the evidence, at the time of the inventions. I understand that such preponderance of the evidence is satisfied if the proposition is more likely to be true than not true.

16. Counsel has informed me, and I understand, that the obviousness inquiry requires that the prior art be considered in its entirety. I am further informed and I understand that an invention cannot be obvious to try where “the breadth of the[] choices and the numerous combinations indicate that the[] disclosures would not have rendered the claimed invention obvious to try.”

17. Counsel has informed me, and I understand, that even where all of the claim limitations are expressly disclosed in the prior art references, there must be some showing that a person of ordinary skill in the art would have been motivated to combine such prior art references and that there would have been a reasonable expectation of successfully achieving the claimed invention from such combination.

18. Counsel has informed me, and I understand, in considering the obviousness of a claimed invention, one should not view the invention and the prior art with the benefit of hindsight. It is for that reason, I am informed and I understand, that obviousness is assessed by the person of ordinary skill in the art at the time the invention was made. In this regard, I am informed and I understand that the invention cannot be used as a guide to selecting and understanding the prior art. I understand that the appropriate standard is to determine whether a person of skill in the art *would* be motivated to combine references, not whether they *could*.

19. Counsel has informed me, and I understand, that obviousness cannot be predicated on what was unknown at the time of the invention, even if the inherency of a certain feature is later established. Counsel has also informed me, and I understand, that unknown properties of the prior art may not be relied upon to provide the rationale for modifying or combining the prior art to reach the claimed subject matter.

20. Counsel has informed me, and I understand, that a reference may be said to teach away when a person of ordinary skill, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the applicant.

21. Counsel has informed me, and I understand, that the “time of invention” applicable to the inventions of claims 6-10 the ‘147 Patent is no later than July 31, 2000, which I understand to be the priority date on the face of the ‘147 Patent.

**C. Person Of Ordinary Skill In The Art**

22. Counsel has informed me, and I understand, that the “person of ordinary skill in the art” (“POSITA”) is a hypothetical person who is presumed to be familiar with the relevant scientific field and its literature at the time of the invention. This hypothetical person is also a person of ordinary creativity capable of understanding the scientific principles applicable to the pertinent field.



23. I am informed by counsel and I understand that the level of ordinary skill in the art may be determined by reference to certain factors, including (1) the type of problems encountered in the art, (2) prior art solutions to those problems, (3) the rapidity with which innovations are made, (4) the sophistication of the technology, and (5) the educational level of active workers in the field. I further understand that the face of the '147 Patent claims a priority date of July 31, 2000.

24. It is my opinion that the person of ordinary skill in the art in the field of the '147 Patent would have been someone with a bachelor's degree in computer science or related field, and either (1) two or more years of industry experience and/or (2) an advanced degree in computer science or related field.

25. Based on my training and experience, I believe that I am (and was as of July 2000) a person of greater than ordinary skill in the relevant art, which permits me to give an opinion about the qualifications of one of ordinary skill at the time of the invention.

26. I note that Dr. Karger's opinion on a person of ordinary skill in the art in his declaration is:

In my opinion, a POSITA at the time of the alleged inventions claimed by the '147 and '069 patents would have a minimum of: (i) a bachelor's degree in computer science, computer engineering, applied mathematics, or a related field of study; and (ii) four or more years of industry experience relating to networking protocols and network

topologies. Additional graduate education could substitute for professional experience, or significant experience in the field could substitute for formal education.

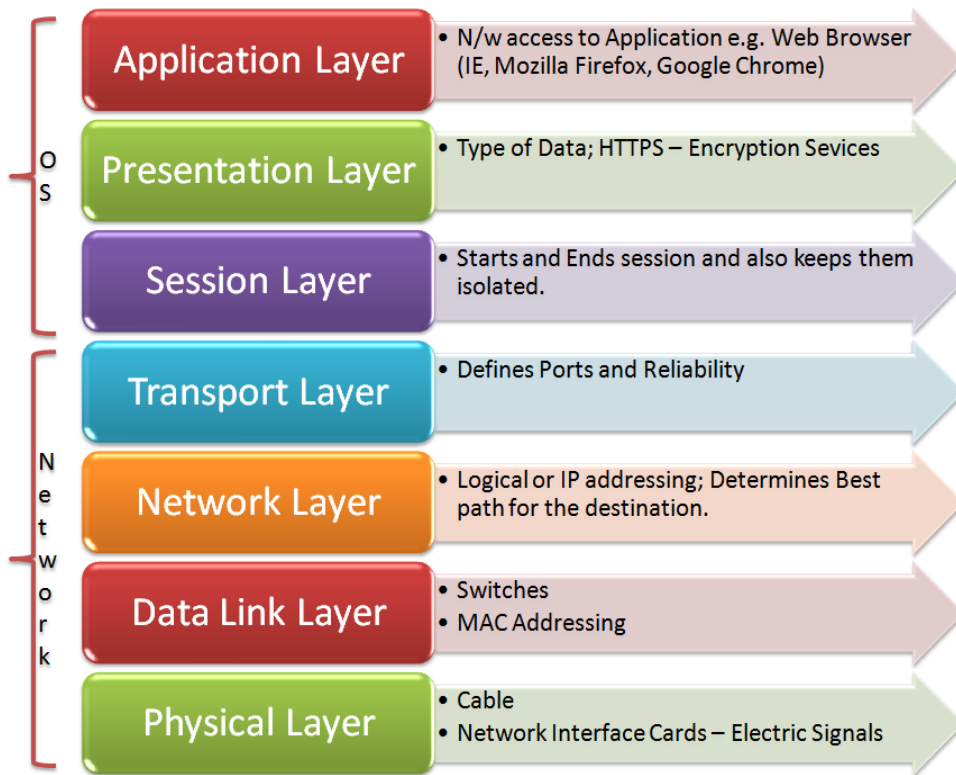
Karger Decl. at ¶ 19.

27. My opinions stated in this declaration would be the same if rendered from the perspective of a person of ordinary skill in the art set out by Dr. Karger.

#### **IV. OVERVIEW OF THE TECHNOLOGY**

##### **A. Network Layers**

28. Network protocols are typically modelled conceptually as being partitioned into layers, which collectively are called the network protocol stack. Each layer provides a set of services and functionality guarantees for higher layers and, to the extent possible, each layer does not depend on details or services from higher levels. To reduce complexity, most networks are designed with a small number of layers, from the physical layer, at the bottom, where computer hardware interfaces with copper wire or wireless radio, to the application layer, at the top, where the user interacts with the software. See Ex. 2007 (“Tanenbaum”) at 17. Each layer uses a set of protocols to build on top of the layers below. The well-known Open System Interconnection (OSI) reference model is used to abstract such layers as shown below:



Ex. 2021; see also Tanenbaum at 28-29.

29. Other protocol stacks also exist, which have overlapping functionality with the OSI Reference Model. The Internet Protocol stack, for example, groups the top three layers as the application layer, while keeping the bottom four layers as in the OSI Reference Model. In the OSI model, the data link layer is used to transfer data between a pair of network nodes or between nodes in a local-area network and to detect errors that occur at the physical layer, dealing with the logical aspects of sending information across network links, such as in the Ethernet protocol. See Tanenbaum at 28-34.

30. The task of the network layer is to provide for the moving of packets between any two hosts, on a best effort basis. It provides a way of individually addressing each host using a numerical label, such as an IP address. This layer is concerned with packet transmission and route searching using routing tables or flooding algorithms. The Internet Protocol (“IP”) is an example of a protocol at the network layer. The task of the transport layer is to support communication based on network (e.g., IP) addresses and ports, which are numerical addresses for higher-level protocols to use. For example, the transport layer in the Internet provides a protocol, the Transmission Control Protocol (TCP), which establishes a virtual connection between a client and server and guarantees delivery of all packets in an ordered fashion, and the Internet also provides a transport protocol, the User Datagram Protocol (UDP), which assumes no prior setup and delivers packets as quickly as possible but with no delivery or ordering guarantees. The role of the session layer is to manage communication sessions, that is, the exchange of information in the form of multiple back-and-forth related transmissions between two nodes, such as in the HTTP protocol, which uses TCP and supports web browsing sessions. See Tanenbaum at 28-34.

31. The presentation layer is dedicated to dealing with the translation of data between a networking service and an application and includes such functions as character encoding, data compression, and encryption and decryption of data, as

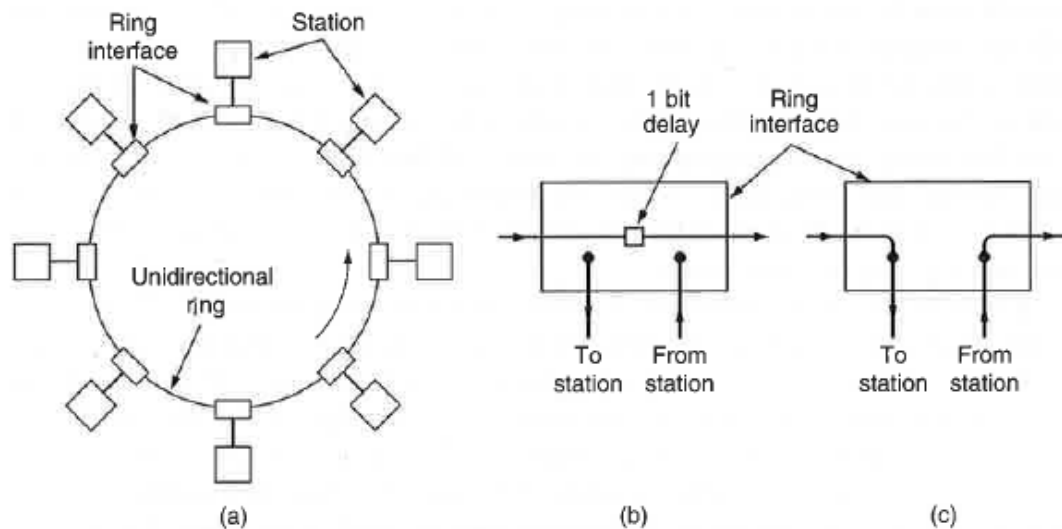
in the HTTPS protocol that is used to encrypt/decrypt private web sessions. The task of the application layer is to provide protocols that support useful functions for users, based on the services provided by the lower layers. Example application-layer functions include web browsers, web-based collaboration systems, and multi-player online games. Typically, the application layer will use application programming interfaces (“APIs”) that will interact with a variety of underlying networks. Designing systems at the application layer is completely different from designing systems at the network layer. An application layer system is typically designed to run on multiple platforms and is usually untethered to the underlying network that may be used to send data to other applications. In contrast, a system at the network layer does not consider the data that is routed through the system or how the application uses the data because its primary job and concern is the getting packets from the source all the way to the destination. See Tanenbaum at 28-34.

## **B. Token Ring Networks**

32. The token ring topology is a simple network topology where stations are connected one to the next so as to form a physical ring. Each station has a single transmitter and a single receiver, such that the transmitter from one station is directly linked to the receiver of another station. *See, e.g.*, Tanenbaum at 292-293. This convention automatically implies that the connections form a circular ring

topology, with messages being transmitted around the ring in the order determined by the transmitters.

33. A special small message, called a “token,” is sent around the stations in such a network in order to facilitate the routing of messages between stations. This token message is continuously transmitted around the ring, in a fashion reminiscent of the classic party game, “hot potato.” That is, a station receives the token from its receiver and transmits it to the next station out its transmitter. The following figure (4-28) from Tanenbaum is an illustration of the topology of a token ring (in part (a) on the left) and the token-passing hardware (in parts (b) and (c) on the right):



**Fig. 4-28.** (a) A ring network. (b) Listen mode. (c) Transmit mode.


Tanenbaum at 293.

34. As referenced in T. Todd, *The Token Grid Network*, IEEE/ACM Transactions on Networking, 2.3, 279–287 (Ex. 1019, “Todd”), the protocol details for sending messages in networks that follow the token ring protocol were standardized in the IEEE 802.5 and FDDI standards. Todd at 279. The token represents a permission to send. For example, according to a typical scenario, in order for a station, X, to send a message to another station, Y, the station X waits until it is holding the token and the token is “empty.” Then it sends the message along with the token to its successor in the token ring. Each time a station, Z, receives the token along with a message, the station Z checks to see if it is the destination for the message. If Z is the destination, then it copies out the message and marks it as sent. In either case, Z transmits the token and message to its successor. When the original sender gets the token and message back again, it checks to see if it was correctly received by the destination and then it empties out the token and passes it to its successor, for the process to enable the possibility of another station to send a message around the ring. Thus, if there are N stations in a token ring, it will take up to N-1 steps to send a message and up to N steps before another message can be sent around the ring.

**C. Token Bus Networks**

35. Another network topology is the token bus topology. In this topology, stations are all connected to a physical bus, which is a broadcast medium such as a

cable, wire, or metal bar capable of conducting electric signals. In a token bus, when any station transmits on the bus, every other station immediately detects the signal. But if two or more stations transmit at the same time, then their signals can interfere with each other and it is likely that no message is sent or received. Thus, the stations need some protocol of managing when they can send messages, with one of the most common being to organize the stations into a logical (not physical) ring and use a small message token to facilitate message passing. The following figure from Jose Rufino *et al.*, *A Study on the Inaccessibility Characteristics of ISO 8802/4 Token-Bus LANs*, IEEE INFOCOM '92: The Conference on Computer Communications (Ex. 1011, "Rufino") is an illustration of this topology, where stations are coupled to the physical bus at the bottom of the figure but are logically organized in a ring with a token message (shown in the figure as a dagger) being passed "around" this logical ring.



Rufino, Figure 1, at 6.



Similarly, the following figure (4-25) from Tanenbaum shows another view of a token bus and its logical ring, including a station that is not currently in the logical ring.

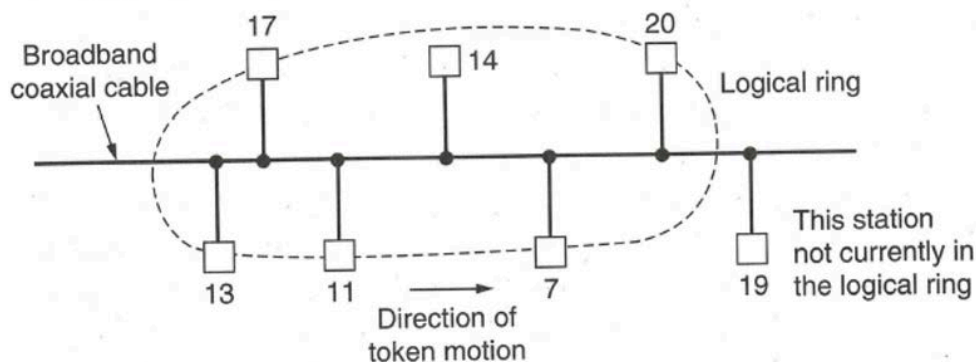


Fig. 4-25. A token bus.

Tanenbaum at 287.

36. As referenced in Rufino, this network topology was standardized in the ISO 8802/4 Token-Bus LAN standard. Rufino at 1. Stations are numbered in a way so that each station knows the identity of its successor station. When a station is holding the token, it is allowed to broadcast a message on the bus, if it has a message ready to send. After broadcasting such a message, or if the station has no message to send, then the station with the token broadcasts the token, identifying its successor as the recipient. Thus, the token “bounces” up and down between stations and the bus in way that logically follows the ordering of stations around the logical ring. In this way, unlike in the token ring topology, a message is sent from a source to a destination in a single step, even if the network has N stations.

Still, a station wishing to send a message might have to wait as long as N-1 steps before it has the token and has permission to use the bus to send its message.


37. There are some important differences between the token ring and token bus topologies. For example, they have different detailed standards for how messages are formulated, how tokens and message are passed, and how networks are setup. For instance, token ring topologies typically follow the IEEE 802.5 or FDDI standards (e.g., see Todd at 279), whereas token buses typically follow the ISO 8802/4 standard (e.g., see Rufino at 1). Given these differences, a POSITA would not find it obvious how to translate a protocol for a token ring into a protocol for a token bus, or vice versa, without detailed instructions on how to modify the detailed protocols and steps used to send messages and perform overhead operations for a token ring into corresponding functionality in a token bus, or vice versa.

38. For example, as I mention above, messages are sent on a broadcast medium in a token bus, whereas they are sent incrementally around the ring in a token ring, resulting in different message-sending performance and protocol details. In addition, as a POSITA would immediately recognize, a station in a token ring can identify its successor simply as the station on the receiving end of its transmitter, whereas a station in a token bus must keep track of the identity of its successor.

39. Indeed, this latter issue is a major factor in the different ways in which stations in the respective two topologies can recover from unplanned station failures. In the case of a token bus, for example, as disclosed in Rufino, several steps are needed to recover from a station failure, including detecting a station failure, having some station respond on the bus for a new successor to be found. *See e.g.*, Rufino at 7-8. A station failure in a token ring, on the other hand, as disclosed in Todd, can be effectively handled with no message passing at all, simply by having a failsafe mechanism whereby a failed station is bypassed so as to link a receiver line for the station to a transmitter line for the station. Todd at 286; *see also* Tanenbaum at 292.

## **V. SUMMARY OF THE ‘147 PATENT**

40. The ‘147 Patent is one of several related patents directed to its novel computer network technology. More particularly, the ‘147 Patent describes methods for disconnecting a node (either in a planned or unplanned manner) from a broadcast channel that overlays a point-to-point network where each node, or participant, is connected to some—but not all—neighboring participants. For example, Fig. 2 of the ‘147 Patent, reproduced below, shows a network of twenty participants, where each participant is connected to four other participants:



**Fig. 2**

*Id.* at Fig. 2.

41. Such a network arrangement, where each node in the network, is connected to the same number of other nodes, is known as an ***m-regular network***.

*Id.* at 4:40-41. That is, a network is *m*-regular when each node is connected to *m* other nodes in a steady state, and a computer would become disconnected from the broadcast channel only if all *m* of the connections to its neighbouring nodes fail.

*Id.* at 4:41-44. In Fig. 2 above, *m*=4 because each node is connected to four other nodes of the network. A network is said to be ***m-connected*** when it would take a failure of *m* computers to divide the graph into disjoint sub-graphs, i.e., separate


broadcast channels. *Id.* at 4:44-47. The '147 Patent also describes a computer network in which the number of network participants,  $N$ , (in Fig. 2, this is twenty) is greater than the number of connections  $m$  to each participant (in Fig. 2, this is four). *Id.* at Fig. 2. This network topology, where no node is connected to every other node, is known as an *incomplete graph*.

42. The incomplete graph topology relies on participants to disseminate information to other participants. *See id.* at 1:60-2:15. As described in the '147 Patent, to broadcast a message, the originating computer sends the message to each of its neighbors using the overlay network. *Id.* at 7:30-35. Each computer that receives the message then sends the message to its neighbors using the network. *Id.* at 7:36-48. In this way, the message is propagated to each computer of the overlay network using the underlying network, thus broadcasting the message to each computer over a logical broadcast channel.

43. The invention claimed in the '147 Patent focuses on processes for removing nodes, or participants, from an existing network or healing disconnections when nodes fail or abruptly leave the network. A computer connected to the network can leave in either a planned or unplanned manner. *Id.* at 8:66-67. The challenged claims at issue relate to unplanned disconnections. For example when a computer leaves the network in an unplanned manner—such as in the event of a power failure—its neighbors discover the computer's absence when

each unsuccessfully attempts to send a message to the disconnected computer. *Id.* at 9:27–31. When a computer detects that one of its neighbors is disconnected, it broadcasts a port connection request over the broadcast channel indicating that it has an open port that needs a connection and identifying the call-in number for the port. *Id.* at 9:33–38. When the port connection request is received at another computer connected to the broadcast channel that has an open port, the other computer contacts the requesting computer to establish a connection. *Id.* at 9:38–42.

44. Fig. 5B of the ‘147 Patent, reproduced below, illustrates the procedure for healing a disconnection of a computer in an unplanned manner. In particular, Fig. 5B shows an exemplary procedure for computer **H**’s neighboring computers to establish new connections when they discover that computer **H** has disconnected in an unplanned manner (*e.g.* without sending a list of its neighboring computers). *Id.* at 9:42–45. When each of computer **H**’s neighbors discovers that **H** has disconnected, it broadcasts a port connection request indicating that it needs to fill an empty port. *Id.* at 9:45–48. In the example shown in Fig. 5B, the dashed lines indicate that computers **F** and **I** respond to each other’s request and form a connection, and computers **A** and **E** respond to each other’s request and form a connection. *Id.* at 9:48–51.



**Fig. 5B**

‘147 Patent, Fig. 5B.

## VI. CLAIM CONSTRUCTION

45. I understand that during *inter partes* review, claims must be “given their broadest reasonable interpretation consistent with the specification.” The broadest reasonable interpretation does not mean the broadest possible interpretation. Rather, the meaning given to a claim term must be consistent with the use of the claim term in the specification and drawings. Thus, even under the broadest reasonable interpretation, the board’s construction cannot be divorced from the specification and the record evidence, and must be consistent with the one that those skilled in the art would reach.

46. In my opinion the person of ordinary skill in the art would adopt the following claim constructions consistent with the broadest reasonable construction of these terms in view of the specification of the ‘147 Patent.

**A. “healing a disconnection of a first computer from a second computer”**

47. In my opinion the person of ordinary skill in the art would understand the term “healing a disconnection of a first computer from a second computer” is properly construed in the context of the ‘147 Patent as “connecting a first computer to another computer in order to maintain an m-regular network after the second computer has involuntarily been disconnected.”

48. The ‘147 Patent discloses that a computer connected to a broadcast channel can leave in a planned manner or in an unplanned manner. ‘147 Patent at 2:59-63, 8:66-9:51; Figs. 5A, 5B. The term “healing a disconnection of a first computer from a second computer” refers to the situation where a computer leaves in an unplanned manner. This is diagrammed in Fig. 5B of the ‘147 Patent. ‘147 Patent at 9:27-51; Fig. 5B. The file history of the ‘147 Patent supports the definition, as well, because it states: “Claims [6-10] relate to the situation where a computer is involuntarily disconnected from the M-regular graph computer network. Claim [6] has been amended to indicate that the healing process of the computer network is performed in a way such as to maintain the M-regular graph nature of the computer network.” Ex. 1002 (Vol. 3) at 690, ‘147 Patent File



History. In my opinion, a POSITA would understand from the specification and the file history that the term, under the BRI standard, “healing a disconnection of a first computer from a second computer” to mean “connecting a first computer to another computer in order to maintain an m-regular network after the second computer has involuntarily been disconnected.”

49. I understand that neither Petitioner nor Dr. Karger provided a definition or construction of “healing a disconnection of a first computer from a second computer”.

**B. “broadcast channel”**

50. In my opinion the person of ordinary skill in the art would understand the term “broadcast channel” is properly construed as “logical channel that overlays an underlying point-to-point communications network.”

51. The specification of the ‘147 Patent supports the construction that the broadcast channel overlays the underlying network. See ‘147 Patent at 4:5-24 (discussing the broadcast channel as an overlay over the network layer), 4:32-36 (describing sending messages using the broadcast channel). Therefore, under the BRI standard, a POSITA would understand that the term “broadcast channel” is a logical channel that overlays an underlying point-to-point communications network.

52. I understand that neither Petitioner nor Dr. Karger provided a definition or construction of “broadcast channel”.

**C. “connected/connection”**

53. In my opinion the person of ordinary skill in the art would understand the term “connected” is properly construed as “having a connection,” and the term “connection” to mean “an edge between two application programs connected to a logical broadcast channel that overlays an underlying network.” This is supported by the specification of the 147 Patent.

54. The specification of the ‘147 Patent explains that the term “connected” refers to the edges connecting participants. In particular, the specification states that “each of the edges represents an ‘edge’ connection between two computers of the broadcast channel.” ‘147 Patent 4:51–53; Fig. 1. Further the specification identifies that the broadcast channel is made of application programs connecting. ‘147 Patent at 15:15-47.

55. Therefore, under the BRI standard, a POSITA would understand that the term “connected” as “having a connection” and the term “connected” as “an edge between two application programs connected to a logical broadcast channel that overlays an underlying network.”

56. I understand that neither Petitioner nor Dr. Karger provided a definition or construction of “port ordering algorithm”.

## VII. SUMMARY OF THE SHOUBRIDGE REFERENCE

57. I reviewed Peter J. Shoubridge et al., *Hybrid Routing in Dynamic Networks*, IEEE International Conference on Communications (Ex. 1005, “Shoubridge”). Shoubridge presents a hybrid routing protocol (via an adaptive routing algorithm) for sending packets at the network layer in the presence of dynamic link failures. Shoubridge at 1381. Shoubridge uses a simulation model because he was unable to test his system in real-life. Shoubridge at 1382.

Shoubridge defines its network model as follows:

“Consider the network model as a directed graph  $G$  with  $N$  nodes and  $L$  bidirectional links. Each node functions as a source of user traffic entering the network where traffic can be destined to all other nodes within the network. Routing algorithms are executed **within these same nodes** and update control messages are exchanged between neighbouring nodes for the purpose of maintaining routing tables.”

Shoubridge at 1382, emphasis added.

58. As part of his simulation model, Shoubridge views the nodes of the network graph as fixed, although the edges can fail and come back up dynamically. In particular, Shoubridge states the following:

“Since the routing procedures are triggered by link cost change events,  $G$  has essentially fixed topology with links always existing between nodes and their respective neighbours. Topology change is modelled by having link quality changing between one of two possible states; **a link is either up providing** good error free

transmission, or a link has failed providing **no useful transmission at all**. This simplified approach requires no topology control as a particular node's neighbours at any given time always belong to the same subset of nodes **as originally defined in G.**”

Shoubridge at 1382, emphasis added

59. Shoubridge describes a hybrid routine strategy, which is described as a blending of a flooding broadcast protocol and a routing-table approach based on routing along shortest paths in the network. With respect to the flooding protocol, Shoubridge states the following:

“Any user packet transmitted from a node is copied and broadcast on all outgoing links. Intermediate transit nodes do not broadcast a packet on the same link that a packet was originally received on. Constrained flooding uniquely identifies packets associated with a particular flood search by using sequence numbering. Nodes store sequence numbers of packets already flooded. If any packets revisit a node with the same sequence number, they are discarded instead of being further broadcast to neighbours. This technique ensures that all nodes are visited at least once and duplicated traffic is kept to a minimum throughout the network.”

Shoubridge at 1382-3.

60. Shoubridge describes a specific network for doing such routing, which is a 64 node network with connectivity of degree 4. Shoubridge states, “[t]he

network is a large regular graph forming a [M]anhattan grid network that has been wrapped around itself as a torus to avoid edge effects.” Shoubridge at 1383.

61. Besides this 8-by-8 configuration, a POSITA would also understand that there is a non-degenerate 4-by-16 configuration, as well. Such topologies are referred to as “Manhattan” networks, because each node is connected to neighbors in the directions North, South, East, and West, much like the streets of Manhattan. One could also imagine a 2-by-32 or 1-by-64 configuration, but a POSITA would understand these to be degenerate, since nodes in such configurations would not have 4 distinct neighbors. A POSITA would understand Shoubridge is referring to non-degenerate configurations here.

62. Neighbor connections in the 64-node degree-4 Manhattan network with torus wrap-around are chosen according to the North, South, East, and West directional connection rule. A POSITA would understand this to be accomplished, for example, by numbering each node with a unique pair of integers,  $(i,j)$ , such that  $i$  and  $j$  range from 0 to 7 and each node  $(i,j)$  is connected to the nodes  $(i-1,j)$ ,  $(i+1,j)$ ,  $(i,j+1)$ , and  $(i,j-1)$ , with all arithmetic done modulo 8, to make the connections according to the North, South, East, and West directions, respectively. Thus, Shoubridge discloses neighbors defined according to a precise mathematical formula or directional relationship, rather than a random selection process.

63. Moreover, Shoubridge is silent on how to add or remove nodes from a Manhattan grid network because Shoubridge designs its simulation such that  $G$  has a fixed topology and that no nodes ever leave or are added. Shoubridge at 1382. Furthermore, a POSITA would understand that there is no simple way to add or remove a node from a 64-node degree-4 Manhattan (8-by-8) grid-with-torus-wrap-around network and have it still be a (non-degenerate) degree-4 Manhattan grid-with-torus-wrap-around network. For example, such a 63-node Manhattan graph would have to be configured as a 7-by-9 or 3-by-21 grid and such a 65-node Manhattan graph would have to be configured as a 5-by-13 grid, both of which would require substantial overhead in converting from an 8-by-8 grid. Moreover, a POSITA would understand that there is no non-degenerate  $n$ -node degree-4 Manhattan grid-with-torus-wrap-around network any time  $n$  is prime. In addition, a POSITA would understand that a 64-node degree-4 Manhattan network with torus wrap-around is 4-connected according to Dr. Karger's definition only if it is in a non-degenerate configuration, such as an 8-by-8 configuration or a 4-by-16 configuration. A 64-node degree-4 Manhattan network with torus wrap-around would not be 4-connected by Dr. Karger's definition, for example, if the network were in a 2-by-32 or 1-by-64 configuration. Therefore, the technical details and requirements of Manhattan grid-with-torus-wrap-around networks make them unsuitable for broadcast channel.

64. Another property of the 64-node degree-4 Manhattan network with torus wrap-around disclosed in Shoubridge is that it is not an overlay network built on top of a point-to-point network, such as the Internet. Instead, Shoubridge discloses a network with explicit connections, such as links in a mobile radio network or wired network. For instance, Shoubridge states with respect to the 64-node degree-4 Manhattan network with torus wrap-around, “[l]ink capacity is fixed at 16kbps to represent a narrowband radio link or perhaps a logical signaling channel.” Shoubridge at 1383. Shoubridge never mentions overlay networks, and instead states, “[f]or efficient routing, nodes require knowledge of network topology.” Shoubridge at 1381. In my opinion, Shoubridge does not teach one skilled in the art how to design or build an m-regular application layer overlay. Shoubridge merely shows an adaptive routing algorithm at the network layer, which is completely different. It would require a POSITA to perform a substantial amount of work to create Shoubridge in a real world network, let alone going from there to an m-regular, incomplete graph application layer overlay.

#### **VIII. SUMMARY OF THE DENES REFERENCE**

65. I also reviewed Tamás Denes, *The “Evolution” of Regular Graphs of Even Order by their Vertices*, Matematikai Lapok, 27, 3–4 (Ex. 1016, Ex. 1017 “Denes”). Of particular relevance in Denes is the  $ET^{-1}$  transformation (Equation 3). Denes at 13.

66. Dr. Karger, in his declaration, explains, using discussions, illustrations, and equations, the  $ET^{-1}$  transformation of Denes, stating, “Denes also discloses that a  $k$ -regular non-complete topology may be maintained when a node is removed from a network.” Ex. 1003 at 76-81. ***In fact, Denes discloses almost the exact opposite of this.*** Dr. Karger is correct in summarizing the actions on a  $k$ -regular graph of the  $ET^{-1}$  equation under the limited circumstances when it can be applied (Denes uses “ $k$ ” instead of “ $m$ ,” but a POSITA would not be confused by this). But he seems to be overlooking crucial details of Definition 3, which precedes this equation, as well as the statement of the equation itself (as well as theorems that follow it in Denes). I reproduce both Definition 3 and the statement of the  $ET^{-1}$  equation below:

Definition 3. Let  $\Gamma_{n,k} = (P, E) \in \mathbf{PR}_k$ ,  $p \in P$ , and we take the vertex pairs  $p_1p_2, p_3p_4, \dots, p_{k-1}p_k$  to be neighbors of  $p$  but pairwise not neighbors. Then we take the  $ET$  transformation to be:

$$(3) \quad ET^{-1}: \Gamma_{n,k} \rightarrow \Gamma_{n-1} = (P', E'), \quad P' = P \setminus \{p\} \\ E' = E \setminus \{(pp_1), (pp_2), \dots, (pp_k)\} \cup \{(p_1p_2), (p_3p_4), \dots, (p_{k-1}p_k)\}$$

Ex. 1016/1017 at 13.

67. A POSITA would understand from these disclosures that the  $ET^{-1}$  equation applies only when both of the following conditions, among others, apply:

- The node  $p$  and all its incident edges must be members of the original graph,  $\Gamma_{n,k}$ . If  $p$  and its incident edges are not members of  $\Gamma_{n,k}$ , then



the equation doesn't make sense, because  $p$  appears in the equation and Definition 3 requires  $p$  to be in the graph.

- There are  $k/2$  vertex pairs,  $p_1p_2, p_3p_4, \dots, p_{k-1}p_k$ , that are neighbors of  $p$  but pairwise not neighbors.

68. Indeed, these conditions are so important that they form the foundations of Denes' Theorem 2, preceded by Definitions 4 and 5, which I reproduce below:

Definition 4. Let  $\Gamma_{n,k} = (P, E)$  be a graph on  $n$  vertices,  $p \in P$ , and  $q_1, q_2, \dots, q_r$  be neighbors of  $p$  in  $\Gamma_n$ . Then we take the subgraph of  $\Gamma_n$  generated by  $p$  (designate this  $G_p = (P', E')$ ), and it turns out this is the subgraph where  $P' = \{q_1, q_2, \dots, q_r\}$  and  $(q_i, q_j) \in E' \Leftrightarrow (q_i, q_j) \in E$ .

Definition 5. Let  $\Gamma_n \in \mathbf{PR}$  and  $p \in \Gamma_n$ . Then the vertex  $p$  is T-characterized if the complement of the subgraph on  $\Gamma_n$  generated by it contains a 1-factor.

Theorem 2. The  $ET^{-1}$  transformation can be applied to a graph  $\Gamma_{n,k} = (P, K) \in \mathbf{PR}_k$  if and only if there exists a T-characterized vertex  $p \in \Gamma_{n,k}$ .

Ex. 1016/1017 at 13.

69. A POSITA would understand that Denes is explaining here that, although some readers might think that the  $ET^{-1}$  equation can always be applied to maintain a graph to be  $m$ -regular after the planned removal of a vertex,  $p$ , this equation actually applies only in limited circumstances, that is, when the graph currently contains the vertex,  $p$ , and when  $p$  is what Denes calls a "T-characterized" vertex, which is inclusive of the bulleted conditions I list above. But Denes does not stop there.

70. Denes goes on to show, for example, in his Theorem 3, that there are an infinite number of graphs that contain no T-characterized vertex. In other words, the *ET<sup>-1</sup> equation cannot be used to remove any of the vertices in such graphs*. I reproduce Theorem 3 from Denes below:

Theorem 3. For every even  $k \geq 4$  there is a graph  $\Gamma_{n,k} \in \mathbf{PR}_k$  which has no vertex  $p$  which is T-characterized.

Ex. 1016/1017 at 13.

71. Moreover, perhaps to show just how common such  $k$ -regular graphs can be, Denes illustrates such a graph (for  $k=4$ ), for which the  $ET^{-1}$  equation cannot be applied to any of its vertices, in his Figure 2, which I reproduce below:




Figure 2

Ex. 1016/1017 at 14.

72. A POSITA would understand that the  $ET^{-1}$  equation cannot be applied to any vertex in this 4-regular graph, because none of the vertices in this graph are T-characterized. In a nutshell, the main problem is that every vertex in this graph has three neighbors that are joined in a triangle, which implies that no vertex has two pairs of neighbors that are themselves not already pairwise neighbors; hence,

the  $ET^{-1}$  equation cannot be used to maintain this graph to be 4-regular after the removal of any of this graph's vertices. Furthermore, a POSITA would understand that having three of one's neighbors form a triangle is a common occurrence in real-world graphs, such as occurs in social networks when someone has three friends who are themselves mutual friends. Thus, a POSITA would understand from reading Denes that *the  $ET^{-1}$  equation does not alone teach a transformation that works in a manner as to maintain an m-regular graph.*

73. In addition, a POSITA would also understand that the  $ET^{-1}$  equation does not apply to an unplanned removal of a vertex, because such a vertex is not a current member of the graph in that circumstance. This might at first seem like a mere technicality, but a POSITA would understand that this is not the case. In an unplanned removal of a vertex, p, the neighbors of p might not know identities of p's other neighbors and, until they try to communicate with p, they might not even know that p is no longer a member of the network. That is, the information related to p and its neighbors is no longer available for trying to repair an m-regular graph after the unplanned removal of p. Thus, a POSITA would understand that repairing an unplanned removal of a vertex from a network requires a completely different set of transformations than the  $ET^{-1}$  transformation. To sum up, if p is not currently a member of the network (as is required by the  $ET^{-1}$  equation), then the

information represented by  $p$  is not available to repair the graph to be  $m$ -regular again.

74. Shoubridge, Rufino, and Todd do not remedy the above deficiencies of Denes, as none of these references disclose the messages and topological changes necessary to heal an unplanned removal of a vertex from an  $m$ -regular network in a manner as to maintain an  $m$ -regular graph. Indeed, my understanding is that Petitioner is pointing only to Denes for teaching this limitation, which, as I have explained, is, in fact, missing from Denes.

75. The '147 Patent, in contrast, explains how to repair an unplanned removal of a vertex (as well as a planned removal) from an  $m$ -regular graph in a manner as to maintain an  $m$ -regular graph. For example, the '147 Patent explains how to heal the situation that arises when two neighboring participants each have an empty internal port and each need to be connected to other participants, which is a situation the '147 Patent calls the “neighbors with empty ports” condition. '147 Patent at 9:52-11:29 and Figures 5C-5F. This technique requires the first computer to recognize the condition to send a “condition check message to the other neighbor[, which] includes a list of the neighbors of the sending computer” to a receiving computer. *Id.* at 9:66–10:4. The receiving computer then sends “a condition repair request to one of the neighbors of the sending computer which is not already a neighbor of the receiving computer[, which then] disconnects from

one of its neighbors... and connects to the computer that sent the condition repair request.” *Id.* at 10:7–14. Finally, the two computers still in need of a connection connect to each other (if they are not already neighbors) or, if they are neighbors, “repeat the condition repair process until two non-neighbors are in need of connection.” *Id.* at 10:16–24. In my opinion, Petitioner’s citation to Denes only goes to show how that there was a long felt need for the invention of the ‘147 Patent, as the ‘147 Patent provides an elegant solution where Denes could not.

#### **IX. SUMMARY OF THE RUFINO REFERENCE**

76. I also reviewed Jose Rufino *et al.*, *A Study on the Inaccessibility Characteristics of ISO 8802/4 Token-Bus LANs*, IEEE INFOCOM ’92: The Conference on Computer Communications (Ex. 1011, “Rufino”).

77. Rufino discloses a number of methods for managing a token bus consistent with the ISO 8802/4 Token-Bus LAN standard, with particular attention paid to inaccessibility characteristics, such as would cause network partitions. Rufino at Abstract. In particular, Rufino states the following:

A study of the behaviour of an ISO 8802/4 token-bus with regard to partitions is the central issue of this paper. Token-bus, given its connection with MAP is very relevant in control and automation<sup>3</sup>, where real-time, reliability and accessibility are a must.

This study contributes to a better understanding of the ISO 8802/4 Token-Bus LAN operation having these attributes in mind.

---

<sup>1</sup> LAN level information packet.

<sup>2</sup> User level information packet.

<sup>3</sup> Manufacturing Automation Protocol.

Rufino at 1, right column (emphasis added).

78. As mentioned elsewhere herein, a token bus topology involves stations that are all connected to a shared broadcast medium, such as a metal cable or bar, and are organized in a logical (not physical) ring, e.g., by having each station store the identity or address of its successor. This topology does not form an m-regular network for m being at least 3, of course.

79. Rufino discusses logical ring “house keeping” functions and actions for the sake of the preservation of the integrity of the token bus network, stating the following:

Several types of protocol data frames are exchanged between peer MAC entities, in order to provide logical ring housekeeping functions, essential for normal ring membership management, and for the preservation of ring integrity in the presence of faults. These functions are supported through the following set of actions:

- Solicit new stations for logical ring entry.
- Find out the successor of a failed station (*who follows*).
- Solicit any station to respond at successor querying.
- Resolve contentions.
- Establish a new successor.
- Bid for token generation.
- Token passing.

Rufino at 3, left column.

80. Of particular relevance in the context of the '147 Patent, are actions disclosed in Rufino for dealing with station failures. For example, Rufino explains that the failure of one station can be detected by the lack of an expected token-passing message. In a variant of a *solicit\_successor* action, the predecessor of the failed station issues a *who\_follows* queries to recover from the station failure. This predecessor station knows it needs to perform this action, because it is the station that just successfully broadcast the token that then was then not resent, and this station knows the token was not resent because it was listening on the shared broadcast medium (i.e., the bus). The *who\_follows* query issued on the bus by this station requests that the station immediately following the failed station in the logical order respond with a *set\_successor* frame containing its identity. Once this

station successfully responds with a *set\_successor* frame, then the logical successor of the issuer of the *who\_follows* query is set in the querying station, and the normal token-passing protocol can resume. Rufino at 5-7.

81. In my opinion, Rufino's methods apply particularly to token bus architectures in the ISO 8802/4 token-bus LAN standard and, not, for example, to token ring architectures, as is outlined elsewhere herein. For example, Rufino states, "Access control is performed by a token passing protocol, ***which establishes a logical ring over the physical bus***. Access to the ***shared broadcast medium*** for data transmission is only granted to the station which currently holds the token." Rufino at 2, right column (emphasis added). In addition, actions such as *solicit\_successor*, *who\_follows*, and *set\_successor* depend critically on the fact that the stations are connected to a physical broadcast medium (i.e., a bus) and these methods would not work or would be unnecessary in a token-ring architecture. For example, if a station fails in token ring network, then it either partitions the network so that messages cannot be sent from one side to the other or it fails with a fallback bypass that connects a receiving line to a transmitting line for that station, which requires no message passing at all.

## **X. SUMMARY OF THE TODD REFERENCE**

82. I also reviewed T. Todd, *The Token Grid Network*, IEEE/ACM Transactions on Networking, 2.3, 279–287 (Ex. 1019, "Todd").



83. Todd describes a token grid network topology, which is a generalization of the token-ring topology to a multiple-dimensional grid. Todd summarizes its contributions as follows:

In this paper, a network referred to as the token grid is introduced. The token grid is a multidimensional extension of the token ring [2], [4] where stations share access to a mesh formed by a set of overlapping rings. Media access is performed by having a station capture one out of a number of tokens circulating within the network. Couplings between the rings are implemented by the user stations in a very simple fashion and under token control. As in a token ring, the network supports the transmission of variable length packets. In the results presented, a dual token ring is used as a basis for performance comparisons with the token grid. This is because unlike other proposed networks, the station transmission hardware requirements for the token grid are identical to that of a dual ring network such as FDDI [4], [3]. Each station in both networks has exactly two transmitters and two receivers. Also, both networks may be implemented with a very small station latency where only a few bits per station need be buffered.

Todd at 279, right column (emphasis added).

84. In addition to Todd explicitly stating that it is extending the token ring topology, and not, for example, a token bus topology, it would be clear to POSITA that Todd is extending a token ring topology, and not a token bus, to a multi-dimensional grid. For example, as Todd states above, each station has exactly two transmitters and two receivers and Todd does not disclose, for example, that stations have connections to one or more buses. *See also* Todd at 282, right

column. In addition, a POSITA would understand that the citations [2], [3], and [4] are to standards, IEEE 802.5 and FDDI, for token rings:

- [2] ANSI/IEEE Standard 802.5, "Token ring access method and physical layer specifications," 1985.
- [3] F.E. Ross, "An Overview of FDDI: The Fiber Distributed Data Interface," *IEEE J. Select. Areas Commun.*, vol. 7, pp. 1043-1051, Sept. 1989.
- [4] American National Standard, "FDDI token ring media access control (MAC)," ANSI X3.139, 1987.

Todd at 287.

85. The token grid of Todd organizes stations in a grid, where the stations in each row and each column are connected to form a token ring. Todd at 279, right column. Todd illustrates an example implementation of this topology in its Figure 1, which I reproduce below:




Fig. 1. The token grid (4 × 4 example).

Todd at 280.

86. As further evidence that Todd is a generalization of the token ring topology, sending and receiving messages in the token grid of Todd is based on stations waiting to receive row or column tokens, or combination row/column tokens, and then transmitting them rather than broadcasting them. For example, Todd gives a summary of the steps for sending and receiving messages as shown below:

PROTOCOL 1:  
BASIC TOKEN GRID PROTOCOL

<u>Column Token <math>j</math> Arrives To Station <math>(i, j)</math> :</u>	
1.	Capture the token.
2.	If a packet is available, transmit onto Column Ring $j$ .
3.	Hold Column Token $j$ . Wait for Row Token $i$ .
4.	Done.
<u>Row Token <math>i</math> Arrives To Station <math>(i, j)</math> :</u>	
1.	Capture the token if a packet is available for Row Ring $i$ or if Column Token $j$ is being held.
2.	Transmit onto Row Ring $i$ if a packet is available.
3.	If the station is not holding Column Token $j$ then release Row Token $i$ and go to 5.
4.	Enter the SR state. Generate and release Row/Column Token $i, j$ .
5.	Done.
<u>Row/Column Token <math>i, j</math> Arrives To Station <math>(i, k)</math> :</u>	
1.	Capture the token if a packet is available for Row/Column Ring $i, j$ or if $k = j$ .
2.	If a packet is available then transmit onto Row/Column Ring $i, j$ .
3.	If $k \neq j$ then release the token and go to 5.
4.	Generate and release Column/Row Token $i, j$ .
5.	Done.
<u>Column/Row Token <math>i, j</math> Arrives To Station <math>(k, j)</math> :</u>	
1.	Capture the token if a packet is available for Column/Row Ring $i, j$ or if $k = i$ .
2.	If a packet is available then transmit onto Column/Row Ring $i, j$ .
3.	If $k \neq i$ then release the token and go to 5.
4.	Enter the DR state. Regenerate and release Column Token $j$ and Row Token $i$ .
5.	Done.

Todd at 281, right column.

87. Of particular relevance to the '147 Patent, Todd discusses the maintenance and reliability of its token grid topology in the presence of station power failures and link failures. Todd at 286. For example, with respect to initialization and maintenance of tokens, Todd states, "Since each station is connected to two independent rings, procedures similar to those in IEEE 802.5 [2] and FDDI [4] may be adopted." Todd at 286. With respect to station failures, Todd states, "As in FDDI [4], it is assumed that in the physical implementation, optical bypass switches are used to maintain ring integrity following a station power failure. When a power failure occurs, the station may be either bypassed in the SR or DR configuration." Todd at 287, right column. A POSITA would understand that Todd does not disclose any message passing as occurring to implement these DR and SR bypass operations, which Todd respectively illustrates in its Figures 2a and 2b, which I reproduce below (where a POSITA would understand that in-coming arrows illustrate receivers and outgoing arrows illustrate transmitters):




Fig. 2. Station configurations.

Todd at 280.

## XI. ANALYSIS OF CLAIMS AT ISSUE

88. I have reviewed the Board's decision to institute trial regarding whether Shoubridge in view of Denes, Rufino, and/or Todd renders obvious certain claims of the '147 Patent. In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious Claims 6-10 of the '147 Patent.

89. Indeed, Shoubridge, Denes, Rufino, and Todd are not relevant to the subject matter of the '147 Patent. A POSITA would not find it obvious to modify

Shoubridge's networking layer system to teach the application layer system of the '147 Patent, nor combine with it Denes, Rufino, and Todd to address the invention. The Board failed to consider a number of important technical distinctions between Shoubridge, Denes, Rufino, and Todd and the '147 Patent, including:

- Shoubridge does not teach a broadcast channel
- Denes does not relate to when a node abruptly leaves a node
- The combination of Rufino and Todd would simply not work because they are discussing two different topics that lends themselves to modification

90. I understand that Dr. Karger failed to consider secondary considerations despite having knowledge of the Sony license. I understand from counsel that secondary considerations are additional factors to be considered in determining obviousness. The factors include commercial success, licensing of the patent, long felt need, copying by others, and other factors. I have reviewed Dr. Harry Bims' report on secondary considerations. In my opinion, the secondary considerations are relevant here, especially because that '147 Patent was licensed to Sony for use in Sony's multi-user gaming platform known as Playstation. The fact that Sony specifically sought to license the patent from Boeing shows that Sony thought the patent was nonobvious and relevant.

91. For each claim limitation, the following analyzes the corresponding portions of Shoubridge in view of Denes, Rufino, and/or Todd cited by Petitioner

to show how a person having skill in the art would understand what the reference actually discloses or teaches.

**A. Claim 6a: “A method for healing a disconnection of a first computer from a second computer, the computers being connected to a broadcast channel, said broadcast channel being an m-regular graph where m is at least 3, the method comprising”**

92. Claim 6a requires a “method for healing a disconnection of a first computer from a second computer, the computers being connected to a broadcast channel, said broadcast channel being an m-regular graph where m is at least 3, the method comprising.” In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

93. For “method for healing a disconnection of a first computer from a second computer, the computers being connected to a broadcast channel, said broadcast channel being an m-regular graph where m is at least 3, the method comprising,” Petitioner cites:

Shoubridge discloses a dynamic computer network including a first computer and a second computer being connected to a broadcast channel, said broadcast channel forming an m-regular graph where m is at least 3. See, e.g., Ex. 1005 at 1381 Abstract (“the distribution of traffic loads and network topologies may vary from nearly static to very dynamic. This dynamic behaviour may vary both in space and in time.”); 1381 ¶ 3 (“For efficient routing, nodes require knowledge of network topology. Links may fail or become congested, nodes or users could be mobile, resulting in changes to source-destination

paths.”); 1383 ¶ 2 (“A 64 node network with connectivity of degree 4 is modeled as  $G$ . The network is a large regular graph forming a Manhattan grid network that has been wrapped around itself as a torus to avoid edge effects.”).

Denes discloses a method of disconnecting a node from another node, the node and the other node being connected in a graph topology, said graph topology forming an  $m$ -regular graph where  $m$  is at least 3. See, e.g., Ex. 1017 at 366 ¶ 4 (“Definition 3. Let  $\Gamma_{n,k} = (P, E) \in PR_k$ ,  $p \in P$ , and we take the vertEx. pairs  $p_1p_2, p_3p_4, \dots, p_{k-1}p_k$  to be neighbors of  $p$  but pairwise not neighbors. Then we take the ET transformation to be: (3) ET-1:  $\Gamma_{n,k} \rightarrow \Gamma_{n-1} = (P', E')$ ,  $P' = P \setminus \{p\}$  and “ $E' = E \setminus \{ (pp_1), (pp_2), \dots, (pp_k) \} \cup \{ (p_1p_2), (p_3p_4), \dots, (p_{k-1}p_k) \}$ ”); 367 Fig. 2.

Todd discloses a token grid network, which “is a two dimensional network structure arranged in  $R$  rows and  $C$  columns. An example for  $R = C = 4$  is shown in Fig. 1.” Ex. 1019 at 1 col. 2 ¶ 4, Fig. 1; Karger ¶ 232. Todd discloses that the reason it devised a token grid topology is because, in 2-regular systems, “the maximum throughput is severely restricted by the data rate of the shared channel.” Ex. 1019 at 1 col. 1 ¶ 2; Karger ¶ 233. And “[i]n principle, this problem may be addressed using a network topology which permits multiple concurrent packet transmissions.” Ex. 1019 at 1 col. 1 ¶ 2; Karger ¶ 233. In other words, Todd proposes adding additional connections in order to allow for better data transmission. Karger ¶ 233.

As discussed in Ground 1, a POSITA would have been motivated to apply the messaging techniques disclosed in Rufino to Shoubridge.



Karger ¶ 234. While a POSITA's would have known how to do so by July 2000, Todd (published in 1994) shows that the modification to Rufino can be done "in a very simple fashion" by overlapping token rings. Ex. 1019 at 1 col. 2 ¶ 2; Karger ¶ 234.

Therefore, Ground 4 simply proposes an alternative ground for independent claims 1, 6, and 11 (which require  $m$  to be at least 3).  
Karger ¶ 235.

Petition at 29, 59-60.

94. Notably Petitioner does not rely on Rufino for this claim element. There is no apparent relation between the portions of Shoubridge, Denes, or Todd cited and method for healing a disconnection of a first computer from a second computer, the computers being connected to a broadcast channel, where the broadcast channel is an  $m$ -regular graph where  $m$  is at least 3. As I discussed above, the proper constructions of "healing a disconnection of a first computer from a second computer," "connected" and "broadcast channel," as understood by a POSITA and as supported by the claims, specification and extrinsic evidence of the '147 Patent, are respectively "connecting a first computer to another computer in order to maintain an  $m$ -regular network after the second computer has involuntarily been disconnected," "an edge between two application programs connected to a logical broadcast channel that overlays an underlying network," "logical channel that overlays an underlying point-to-point communications

network.” A POSITA would understand that a “broadcast channel” under the ‘147 Patent operates at the application layer.

95. In particular, Shoubridge does not teach or render obvious a “method for healing a disconnection of a first computer from a second computer” because Shoubridge does address a situation where a node leaves a network abruptly and the network heals itself. Shoubridge does not teach a network where computers leave and enter an m-regular network. Indeed, as discussed above, Shoubridge teaches a simulation where nodes do not leave because Shoubridge relies on a fixed topology where the link value costs are changed. Shoubridge at 1382 (“*G* has essentially fixed topology with links always existing between nodes and their respective neighbors.”). The choice of never having a new node enter or an old node leave in Shoubridge was, of course, by design to make the simulation simpler. Shoubridge at 1382 (“This simplified approach requires no topology control as a particular node’s at any given time always belong to the **same subset of nodes** as originally defined in *G*.”) (emphasis added).

96. Nor does Shoubridge discuss a situation where it has to heal the network. At best, Dr. Karger cites to generalized statements about links failing or being congested. Ex. 1003 at ¶¶128-131,162. However, that situation is not equivalent to the claims, because Shoubridge never has to repair edges or heal the network, the links are simply either there or not. Shoubridge at 1382 (“Topology

change is modelled by having link quality changing between one of two possible states; a link is either up providing good error free transmission, or a link has failed providing no useful transmission at all.”). There cannot be any repairs as described in the ‘147 Patent in the Shoubridge simulation model, because Shoubridge simulation model does not modify links or connections. See generally, ‘147 Patent at 9:52-10:65; Shoubridge at 1382-1383. This makes sense because the entire purpose of Shoubridge is to discuss adaptive hybrid routing algorithm that is only concerned with routing within the network as it is, not trying to fix connections. Shoubridge at 1381 (discussing routing algorithms). As a result, Shoubridge does not discuss or disclose a situation where nodes leave abruptly, nor how the network heals itself, because the situation does not even arise in Shoubridge.

97. To the extent that Petitioner changes its position, and argues that it would have obvious to a POSITA when applying Shoubridge in the real-world, I disagree. As explained above, Shoubridge has many obstacles to being applied in a real-world manner. Moreover, a real-world Shoubridge would not seek to heal links as, explained above, since its only concern was with routing to every node available not to repair broken links. As such, Shoubridge fails to render obvious this claim element.

98. Moreover, Shoubridge does not teach or disclose a broadcast channel, because it does not discuss an overlay network. Shoubridge is silent with regard to

overlay networks, and instead teaches only an underlying network which routes packets at the network layer. Indeed, Dr. Karger admits that Shoubridge does not teach an overlay network. Ex. 2010, 11/29/16 Karger Tr. 20:17-23 (admitting that the broadcast channel of the '147 Patent is an overlay); *id.* at 48:17-23 (admitting that Shoubridge does not teach an overlay network); *see also* Ex. 2011, 7/6/16 Karger Tr. at 102:22-103:4, Ex. 2012, 7/7/16 Karger Tr. at 81:7-19. Accordingly, it is my opinion that Shoubridge does not render obvious the claim element because it does disclose or render obvious a broadcast channel.

99. It would not have been obvious to apply Shoubridge over the Internet because Shoubridge rewrites and modifies standard routing protocols and is designed to replace the Internet Protocol. Indeed, applying Shoubridge to the Internet is impracticable as flooding does not work in large networks. Ex. 2013, Vu at 624-626, Figs. 3, 4 and 5 (note that the Vu reference is cited by Shoubridge).

100. Specifically, one would not be motivated to create an overlay network with Shoubridge because flooding would have created significant latency and would have prevented a useful dissemination of information through the use of flooding because of the congestion at the higher layer protocols. Ex. 1005 (“Shoubridge”) at 1381 (“flooding is very wasteful of network resources”), at 1384 (“Unfortunately, flooding results in low network utilization and therefore better routing strategies need to be sought.”); Ex. 2013, Vu at 624 (“The abundance of

packets not only puts serious strains on the link capacity, induces higher levels of network congestion, but also dictates a high packet processing capacity at each relay node.”); *see also* Ex. 2007 (“Tanenbaum”) at 351 (describing limited uses of flood search algorithms). Putting Shoubridge’s routing algorithm on the Internet ignores the fact that it’s not a simple design choice, because it requires an understanding of the operation of every layer of the OSI model and the use of standards. A POSITA would not look to Shoubridge and think to make an overlay layer the application level because Shoubridge was specifically about routing algorithms at the network layer. To modify Shoubridge in such a manner is to undermine the point of Shoubridge and is not practical for the reasons stated above.

101. Petitioner also relies on Denes for the element “healing a disconnection of a first computer from a second computer.” Petition at 29-30. In particular, Petitioner cites the formula “ $ET^{-1}: \Gamma^{n,k} \rightarrow \Gamma_n^{-1} = (P', E'), P' = P \setminus \{p\}$ ” and “ $E' = E \setminus \{ (pp1), (pp2), \dots, (ppk) \cup \{ (p1p2), (p3p4), \dots, (pk-1pk) \}$ .” Denes at 366. The “ $ET^{-1}$ ” denotes that the graph is removing a node. Petitioner either failed to recognize or failed to disclose that the formula above only applies in the situation where the leaving node leaves in a planned manner, because, as Denes articulates in prose and in mathematics, the leaving the node, “p,” has to be a member of the graph, “ $\Gamma_{n,k}$ ,” prior to removal. See Denes at 366 (“The  $ET^{-1}$

transformation can be applied to a graph  $\Gamma_{n,k} = (P,K) \in PR_k$  **if and only if** there exists a T-characterized vertex  $p \in \Gamma_{n,k}$ .”) (emphasis added); see also ‘147 Patent at Fig. 5A (graph of orderly removal of node where H is a member of the graph); 5B (graph of unplanned removal of node where H is not a member of the graph).

102. That is, the  $ET^{-1}$  transformation requires explicit knowledge of the leaving vertex,  $p$ . For without  $p$ , the  $ET^{-1}$  equation makes no sense. Thus, the equation for the  $ET^{-1}$  transformation cannot apply to the case, as required in claims 6-10, of a healing of a disconnection of a first computer from a second computer. For to perform a healing of a disconnection of a first computer from a second computer, a POSITA would understand that the second computer is no longer connected to the first computer after limitation 6c of claim 6, because if the second computer were still connected to the first computer, the first computer would not need a connection, as is required in limitation 6c. Moreover, given the temporal language of claim 6, a POSITA would understand that the limitations of claim 6 occur in temporal order, first 6b (attempting to send a message from the first computer to the second computer), then 6c (when the attempt to send the message is unsuccessful, broadcasting from the first computer a connection port search message indicating that the first computer needs a connection), and then 6d (having a third computer not already connected to said first computer respond to said connection port search message in a manner as to maintain an m-regular graph).

Therefore, limitation 6d requires a network transformation that does not involve the second computer; hence, the equation for the  $ET^{-1}$  transformation does not apply to limitation 6d. Accordingly, Denes fails to disclose “healing a disconnection of a first computer from a second computer” or any of the other elements of the claim, and Petitioner has not provided sufficient support as why POSITA would modify Denes to disclose the claimed element.

103. Petitioner also relies on Todd, but only for the limitation of “where  $m$  is at least 3.” However, Petitioner fails to connect this citation with Shoubridge and Denes and provide a rational motivation to combine. Petitioner explicitly does not refer to or cite to Rufino in element 6a. See Petition at 29-30. Petitioner’s arguments regarding Todd only relate to Rufino, but the noticeable absence of Rufino undermines Petitioner’s position.

104. While Rufino is not relied on by Petitioner for this element (and contradictory to Petitioner above cited reliance on Shoubridge), Petitioner does rely on Rufino for “broadcast channel” in the rest of the claim elements without explaining what in Rufino is the broadcast channel. See Petition at 32 (“[T]he station broadcasts a connection port search message on the broadcast channel....”); *id.* at 33 (“[T]he station sends a who follows query (connection port search message) on the broadcast channel....”); *id.* at 36 (“Rufino discloses broadcasting messages on the broadcast channel to maintain a regular graph....”). To the extent

that Petitioner is arguing the physical or the logical token ring layers are broadcast channels, I disagree for the reasons set forth below.

105. I disagree with both theories. First, Rufino's physical bus is not a logical channel, nor does it overlay any communications network, let alone an underlying point-to-point communications network because it simply provides physical communications links between computers as shown in the figure below.




Fig. 4-25. A token bus.

Tanenbaum at 287. As described in my discussion of the token bus network, a special data frame, called a token, “propagates around the logical ring, with only the token holder being permitted to transmit frames. Since only one station at a time holds the token, collisions do not occur” Tanenbaum at 288.

106. Second, the logical token ring is not a broadcast channel because it does not support the broadcasting of messages, but rather only sends “token frames specifically address to its logical neighbour in the ring.” Tanenbaum at 288. The design of the token bus network (IEEE 802.4 Standard) is to have all other frames



used in token bus network sent either directly to another station or sent to a group of stations on the physical bus. Note that ISO 8802/4 discussed in Rufino is the same as IEEE 802.4 Standard. Tanenbaum at 275. The “who follows” query cited by Petitioner, is actually sent on the physical bus. Tanenbaum at 287-288, 290; *see also* Ex. 2010, 11/29/16 Karger Tr. 68:6-9 (“Q. Is the “who follows” query sent on the physical bus? A. Yes. There is a “who follows” query that is transmitted on the physical bus.”).

107. Additionally, the logical token ring does not overlay a point-to-point communications network as required. Instead, the logical token ring overlays a physical medium. *See* Tanenbaum at 287; Rufino at 9 (“Access control is performed by a token passing protocol, which establishes a logical ring over the physical bus.”). Moreover, as discussed elsewhere herein, Rufino logical ring over a physical bus is at best 2-regular and is not extensible to be m-regular where m is at least 3.

108. As explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion, Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

**B. Claim 6b: “attempting to send a message from the first computer to the second computer; and”**

109. Claim 6b requires “attempting to send a message from the first computer to the second computer.” In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

110. For this claim limitation, Petitioner cites:

Shoubridge discloses a dynamic network forming a broadcast channel, said broadcast channel forming an m-regular graph where m is at least 3. See, e.g., Ex. 1005 at 1381 Abstract; 1381 ¶ 3; 1383 ¶ 2.

Denes discloses a method of disconnecting a first node from the second node. See, e.g., Ex. 1017 at 366 ¶ 4; 367 Fig. 2.

Rufino discloses attempting to send a message from the first computer to the second computer. See, e.g., Ex. 1011 at 0958 ¶ 4 (“Uncontrolled omissions and partitions are a source of asynchrony and inconsistency. This is unacceptable for most systems, let alone real-time ones.”); 0959 col. 1 - ¶ 7 – col. 2 ¶ 1 (“Access control is performed by a token passing protocol, which establishes a logical ring over the physical bus.”); 962 col. 2 ¶¶ 4-5 (“Station withdrawal is achieved through a ring patch between its predecessor and successor stations. For that purpose, the leaving station before passing the token issues a set successor frame, addressed to its predecessor, and carrying the address of its future successor.”); 0962 col. 2 ¶ 9 – 0963 col. 1 ¶ 1 (“If this station is failed, the token passing operation will not succeed. After the token pass checking period (one slot time) has elapsed a recovery strategy is tried. This includes the repetition of token

transmission (checked during one more slot time) followed by a variant of the solicit successor procedure. During this who follows query the current token holder waits for a set successor frame, until the end of a three-slot-time period. Under a single station failure assumption, the current token holder will receive, within this period, a response from the successor of the failed station. This establishes a new successor for the current token holder and ends the recovery procedures.”).

Petition at 30-31.

111. There is no apparent relation between the cited portions of Shoubridge, Denes, Rufino, or Todd and the claim element attempting to send a message from the first computer to the second computer. As discussed above, Shoubridge in view of Denes, Rufino, and Todd fails to disclose or render obvious Claim 6(a), thus the combination fails to disclose or render obvious Claim 6(b).

112. In addition, as explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion, Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

**C. Claim 6c: “when the attempt to send the message is unsuccessful, broadcasting from the first computer a connection port search message indicating that the first computer needs a connection”**

113. Claim 6c requires “when the attempt to send the message is unsuccessful, broadcasting from the first computer a connection port search

message indicating that the first computer needs a connection.” In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

114. For this claim limitation, Petitioner cites:

Shoubridge discloses a dynamic network forming a broadcast channel, said broadcast channel forming an m-regular graph where m is at least 3. See, e.g., Ex. 1005 at 1381 Abstract; 1381 ¶ 3; 1383 ¶ 2.

Denes discloses a method of disconnecting a first node from the second node. See, e.g., Ex. 1017 at 366 ¶ 4; 367 Fig 2.

Rufino discloses when the attempt to send a message is unsuccessful, broadcasting from the first computer a connection port search message indicating that the first computer needs a connection. See, e.g., Ex. 1011 at 0962 col. 2 ¶¶ 4-5 (“Station withdrawal is achieved through a ring patch between its predecessor and successor stations. For that purpose, the leaving station before passing the token issues a set successor frame, addressed to its predecessor, and carrying the address of its future successor.”); 0962 col. 2 ¶ 9 – 0963 col. 1 ¶ 1 (“If this station is failed, the token passing operation will not succeed. After the token pass checking period (one slot time) has elapsed a recovery strategy is tried. This includes the repetition of token transmission (checked during one more slot time) followed by a variant of the solicit successor procedure. During this who follows query the current token holder waits for a set successor frame, until the end of a three-slot-time period. Under a single station failure assumption, the current token holder will receive, within this period, a

response from the successor of the failed station. This establishes a new successor for the current token holder and ends the recovery procedures.”).

Petition at 33-34.

115. There is no apparent relation between the cited portions of Shoubridge, Denes, Rufino, or Todd and the claim element when the attempt to send the message is unsuccessful, broadcasting from the first computer a connection port search message indicating that the first computer needs a connection. As discussed above, Shoubridge in view of Denes, Rufino, and Todd fails to disclose or render obvious Claim 6(a), thus the combination fails to disclose or render obvious Claim 6(c).

116. In addition, as explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion, Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

**D. Claim 6d: “having a third computer not already connected to said first computer respond to said connection port search message in a manner as to maintain an m-regular graph”**

117. Claim 6d requires “having a third computer not already connected to said first computer respond to said connection port search message in a manner as to maintain an m-regular graph.” In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

118. For this claim limitation, Petitioner cites:

Shoubridge discloses a dynamic network forming a broadcast channel, said broadcast channel forming an m-regular graph where m is at least 3. See, e.g., Ex. 1005 at 1381 Abstract; 1381 ¶ 3; 1383 ¶ 2.

Denes discloses having a third node not already connected to the first node respond in a manner as to maintain an m-regular graph. See, e.g., Ex. 1017 at 3677 ¶ 4 (“Definition 3. Let  $\Gamma_{n,k} = (P, E) \in PR_k$ ,  $p \in P$ , and we take the vertex pairs  $p_1p_2, p_3p_4, \dots, p_{k-1}p_k$  to be neighbors of  $p$  but pairwise not neighbors. Then we take the ET transformation to be: (3) ET-1:  $\Gamma_{n,k} \rightarrow \Gamma_{n-1} = (P', E')$ ,  $P' = P \setminus \{p\}$ ” and “ $E' = E \setminus \{(pp_1), (pp_2), \dots, (pp_k) \cup \{(p_1p_2), (p_3p_4), \dots, (p_{k-1}p_k)\}$ ”); 367 Fig. 2.

Rufino discloses having a third computer not already connected to said first computer respond to said connection port search message in a manner as to maintain an m-regular graph. See, e.g., Ex. 1011 at 0960 col. 1 ¶ 3-4 (“Several types of protocol data frames are exchanged between peer MAC entities, in order to provide logical ring housekeeping functions, essential for normal ring membership management, and for the preservation of ring integrity in the presence of faults.”); 0962 col. 2 ¶ 5 (“Station withdrawal is achieved through a ring patch between its predecessor and successor stations. For that purpose, the leaving station before passing the token issues a set successor frame, addressed to its predecessor, and carrying the address of its future successor.”); 0962 col. 2 ¶ 9 – 0963 col. 1 ¶ 1 (“If this station is failed, the token passing operation will not succeed. After the token pass checking period (one slot time) has elapsed a

recovery strategy is tried. This includes the repetition of token transmission (checked during one more slot time) followed by a variant of the solicit successor procedure. During this who follows query the current token holder waits for a set successor frame, until the end of a three-slot-time period. Under a single station failure assumption, the current token holder will receive, within this period, a response from the successor of the failed station. This establishes a new successor for the current token holder and ends the recovery procedures.”).

Petition at 34-35.

119. There is no apparent relation between the cited portions of Shoubridge, Denes, Rufino, or Todd and the claim element having a third computer not already connected to said first computer respond to said connection port search message in a manner as to maintain an m-regular graph. As discussed above, Shoubridge in view of Denes, Rufino, and Todd fails to disclose or render obvious Claim 6(a), thus the combination fails to disclose or render obvious Claim 6(d). Moreover, Petitioner only relies on Denes and Rufino for disclosing this claim element. *See* Petition at 34–35; Ex. 1003 at ¶172. However, neither references discloses the claimed element.

120. Petitioner states: “Denes discloses having a third node not already connected to the first node respond in a manner as to maintain an m-regular graph.” Petition at 35. However, this argument is false as Denes does not disclose

any messaging that would permit a node to “respond” to messages. Indeed, Dr. Karger testified that Denes does not send messages between nodes of graphs. *See* Ex. 2010, 11/29/16 Karger Tr. 55:13-16 (“Q. Does Denes disclose sending messages between nodes of graphs? A. I don't recall him doing so. I doubt it.”). Moreover, as discussed previously, the equation relied on by Petitioner only applies to the situation where the node is part of the graph at the time the transformation is applied as opposed to the abrupt leaving scenario recited in claims 6–10. Furthermore, as explained elsewhere herein, the  $ET^{-1}$  equation of Denes alone is not effective to maintain an m-regular graph subject to node removals. As such, this element is not disclosed or rendered obvious by Denes.

121. Petitioner also argues that Rufino discloses the claimed element. However, as discussed above, Rufino cannot meet this claim element because of the design of the physical and logical topologies of Rufino's token bus LAN. As discussed above, and admitted by Dr. Karger, the “who follows” query is sent on the physical bus. As shown in the diagram below, the station that responds to the “who follows” query is already connected to the physical bus.






Fig. 4-25. A token bus.

Tanenbaum at 287.

122. Accordingly, there is no way that this claim element be met by Rufino because that alleged third station is already connected to the alleged first computer.

Tanenbaum explains that fact:

An important point to realize is that the physical order in which the stations connected to the cable is not important. Since the cable is inherently a broadcast medium, *each station receives the frame*, discarding those not addressed to it. When a station passes the token, it sends a token frame specifically addressed to its logical neighbor in the ring, irrespective of where that station is physically located on the cable.

Tanenbaum at 288 (emphasis added).

123. This is by design and regulated by an IEEE Standard (802.4)/ ISO 8802/4. A POSITA would not seek to modify the IEEE Standard (802.4)/ ISO 8802/4 to make a complete redesign of the protocol standard as suggested by Petitioner.

124. In addition, as explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion, Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

**E. Claim 7: “The method of claim 6 including: when a third computer receives the connection port search message and the third computer also needs a connection, sending a message from the third computer to the first computer proposing that the first computer and third computer connect.”**

125. Claim 7 requires “when a third computer receives the connection port search message and the third computer also needs a connection, sending a message from the third computer to the first computer proposing that the first computer and third computer connect.” In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

126. For this claim limitation, Petitioner cites:

Rufino discloses when a third computer receives the connection port search message and the third computer also needs a connection, sending a message from the third computer to the first computer proposing that the first computer and third computer connect. See, e.g., Ex. 1011 at 0962 col. 2 ¶ 5 (“Station withdrawal is achieved through a ring patch between its predecessor and successor stations. For that purpose, the leaving station before passing the token issues a set successor frame, addressed to its predecessor, and carrying the address of its future successor.”); 0962 col. 2 ¶ 9 – 0963 col. 1 ¶ 1 (“If

this station is failed, the token passing operation will not succeed. After the token pass checking period (one slot time) has elapsed a recovery strategy is tried. This includes the repetition of token transmission (checked during one more slot time) followed by a variant of the solicit successor procedure. During this who follows query the current token holder waits for a set successor frame, until the end of a three-slot-time period. Under a single station failure assumption, the current token holder will receive, within this period, a response from the successor of the failed station. This establishes a new successor for the current token holder and ends the recovery procedures.”).

Petition at 37-38.

127. There is no apparent relation between the cited portions of Shoubridge, Denes, Rufino, or Todd and the claim element when a third computer receives the connection port search message and the third computer also needs a connection, sending a message from the third computer to the first computer proposing that the first computer and third computer connect. As discussed above, Shoubridge in view of Denes, Rufino, and Todd fails to disclose or render obvious Claim 6, thus the combination fails to disclose or render obvious Claim 7.

128. Petitioner relies solely on Rufino for this claim element. Petition at 37-38; Ex. 1003 at ¶179. Petitioner also identified that the set\_successor frame is sent on a broadcast channel by the third computer to the first computer. However, set\_successor is sent on the physical bus. *See* Tanenbaum at 287-288, 290. As

discussed above, the physical bus of the token bus network is not a broadcast channel. Accordingly, Petitioner's arguments fail to disclose this claim element. Petitioner also fails provide evidence that POSITA would have been motivated to modify the IEEE standard (802.4)/ ISO 8802/4 to change how "set\_successor" would have operated.

129. In addition, as explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion, Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

**F. Claim 8: "The method of claim 7 including: when the first computer receives the message proposing that the first computer and third computer connect, sending from the first computer to the third computer a message indicating that the first computer accepts the proposal to connect the first computer to the third computer"**

130. Claim 8 requires "when the first computer receives the message proposing that the first computer and third computer connect, sending from the first computer to the third computer a message indicating that the first computer accepts the proposal to connect the first computer to the third computer." In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

131. For this claim limitation, Petitioner cites:

Rufino discloses when the first computer receives the message proposing that the first computer and third computer connect, sending from the first computer to the third computer a message indicating that the first computer accepts the proposal to connect the first computer to the third computer. See, e.g., Ex. 1011 at 0962 col. 2 ¶ 9 – 0963 col. 1 ¶ 1 (“If this station is failed, the token passing operation will not succeed. After the token pass checking period (one slot time) has elapsed a recovery strategy is tried. This includes the repetition of token transmission (checked during one more slot time) followed by a variant of the solicit successor procedure. During this who follows query the current token holder waits for a set successor frame, until the end of a three-slot-time period. Under a single station failure assumption, the current token holder will receive, within this period, a response from the successor of the failed station. This establishes a new successor for the current token holder and ends the recovery procedures.”).

Petition at 39.

132. There is no apparent relation between the cited portions of Shoubridge, Denes, Rufino, or Todd and the claim element when the first computer receives the message proposing that the first computer and third computer connect, sending from the first computer to the third computer a message indicating that the first computer accepts the proposal to connect the first computer to the third computer. Petitioner only relies on Rufino to disclose this element. See Petition at 39; Karger Decl. at ¶¶183-186. As discussed above, Shoubridge in view of Denes,

Rufino, and Todd fails to disclose or render obvious Claim 6 or 7, thus the combination fails to disclose or render obvious Claim 8.

133. Moreover, Petitioner's citation to Rufino fails to disclose the claim element. In fact, Rufino clearly states that there is no communications from the alleged first computer to the alleged third computer accepting the proposed connection. See Rufino at 962 col. 2 ¶ 9-961 col. 1 ¶ 1 ("the current token holder [i.e., alleged first computer] will receive ... a response from the successor of the failed station [i.e., alleged third computer]. This establishes a new successor for the current token holder and **ends the recovery procedure.**") (emphasis added). In other words, there is no message from the alleged first computer to the alleged third computer accepting the proposed connection because no message exists. As explained above, the IEEE standard (802.4)/ ISO 8802/4 sets the standard for this protocol and has designed it as such that there is no message sent because the next station is already connected on the physical bus. Petitioner and Dr. Karger fail to provide any explanation or point to the alleged message they are referring to because the message is non-existent. See Petition at 39; Karger Decl. at ¶¶ 183-186. Accordingly, Petitioner has failed to explain how this element is rendered obvious in view of Shoubridge in view of Denes, Rufino, and Todd.

134. In addition, as explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion,

Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

**G. Claim 9: “The method of claim 6 wherein each computer connected to the broadcast channel is connected to at least three other computers.”**

135. Claim 9 requires “wherein each computer connected to the broadcast channel is connected to at least three other computers.” In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

136. For this claim limitation, Petitioner cites:

Shoubridge discloses wherein each computer connected to the broadcast channel is connected to at least three other computers. See, e.g., Ex. 1005 at 1383 ¶ 2 (“A 64 node network with connectivity of degree 4 is modeled as G. The network is a large regular graph forming a Manhattan grid network that has been wrapped around itself as a torus to avoid edge effects.”).

Petition at 40.

137. There is no apparent relation between the cited portions of Shoubridge, Denes, Rufino, or Todd and the claim element wherein each computer connected to the broadcast channel is connected to at least three other computers. Petitioner relies solely on Shoubridge for this claim. Notably, the identified broadcast channel is different from the one Petitioner relies on for claim elements

6(b)-(d). In addition, as discussed above in Claim 6(a), Shoubridge does not disclose a broadcast channel as required by this claim. Similarly, because Shoubridge in view of Denes, Rufino, and Todd fails to disclose or render obvious Claim 6, the combination also fails to disclose or render obvious Claim 9.

138. In addition, as explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion, Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

**H. Claim 10: “The method of claim 6 wherein the broadcasting includes sending the message to each computer to which the first computer is connected”**

139. Claim 10 requires “wherein the broadcasting includes sending the message to each computer to which the first computer is connected.” In my opinion, Shoubridge in view of Denes, Rufino, and/or Todd does not render obvious this claim limitation.

140. For this claim limitation, Petitioner cites:

Shoubridge discloses wherein the broadcasting includes sending the message to each computer to which the first computer is connected. See, e.g., Ex. 1005 at 1381 ¶ 1 (“In addition to successfully forwarding user traffic between source and destination nodes in a communications network, the routing function generally seek to optimize some performance criteria such as network utilization or throughput, or perhaps average delay.”); 1382 ¶ 11 to 1383 ¶ 1 (“Any



user packet transmitted from a node is copied and broadcast on all outgoing links. Intermediate transit nodes do not broadcast a packet on the same link that a packet was originally received on.”).

Petition at 41.

141. There is no apparent relation between the cited portions of Shoubridge, Denes, Rufino, or Todd and the claim element wherein each computer connected to the broadcast channel is connected to at least three other computers. Petitioner relies solely on Shoubridge for this claim. Notably, the identified broadcast channel is different from the one Petitioner relies on for claim elements 6(b)-(d). In addition, as discussed above in Claim 6(a), Shoubridge does not disclose a broadcast channel as required by this claim. Similarly, because Shoubridge in view of Denes, Rufino, and Todd fails to disclose or render obvious Claim 6, the combination also fails to disclose or render obvious Claim 10.

142. In addition, as explained further below, there is no motivation to combine Shoubridge, Denes, Rufino, and Todd. As such, in my opinion, Shoubridge in view of Denes, Rufino, and Todd fail to render obvious this claim element.

## **XII. THERE IS NO MOTIVATION TO COMBINE SHOUBRIDGE AND DENES AND RUFINO AND TODD**

143. In my opinion POSITA would not be motivated to combine Shoubridge, Denes, Rufino, and Todd. Petitioner argues that POSITA would be

motivated to combine Shoubridge with Denes, Rufino, and Todd because they “address the problem of dynamic networks” and “the problem of maintaining a non-complete regular topology.” Petition at 17-19, 58; Karger Decl. at ¶¶ 119-127. I disagree, because Shoubridge, Denes and Rufino, and Todd are disparate systems that one of skill in the art would not have been motivated to combine.

144. I disagree with Petitioners and Dr. Karger’s assessment. As an initial matter, Shoubridge discloses a hybrid routing algorithm using a simulation model, not a real dynamic network. Shoubridge at 1381. Shoubridge’s “dynamic network” simulates changes that “are evenly distributed across all links.” Shoubridge at 1383. A POSITA would understand this assumption is not realistic in the real world, where links are different geographically and in terms of their implementations. In my opinion, a POSITA would not have not have been motivated to combine Shoubridge with Denes and Rufino because neither Denes or Rufino provide any basis to improve the hybrid routing algorithm, which is the essence of Shoubridge, or to his simulation. Moreover, in Shoubridge uses a fixed topology in which nodes never enter or leave the alleged network. Shoubridge at 1382 (“Since the routing procedures are triggered by link cost change events, **G has essentially fixed topology with links always existing between nodes and their respective neighbours.** ... This simplified approach requires no topology control as a particular node's neighbours at any given time always belong to the

same subset of nodes **as originally defined in G.**”(emphasis added). Given Shoubridge’s fixed node topology, a POSITA would have had no reason to modify Shoubridge with Denes, which describes techniques for adding nodes and edges to and removing nodes and edges from a graph. That is, because *nodes* never fail in Shoubridge’s network—only links between nodes fail—the solution for fixing a “link that has no useful transmission at all”—and maintaining an m-regular network—is simply to restore the link or ignore it. *See* Shoubridge at 1382.

145. Petitioners argue that Denes discloses a mathematical principal to remove a node from a network. Petition at 18. This is irrelevant, however, because it does not address the issue of an unplanned or abrupt leaving of the network by a node. Indeed, as discussed above, Denes’  $ET^{-1}$  transformation equation has only limited use and contains many restrictions. Shoubridge discloses his method of simulating a network node change by adjusting the link costs. Shoubridge at 1383 (“To reflect some stochastic behaviour in the simulation model, link cost change events are scheduled with event times derived from two negative exponential probability distributions.”). The argument that one could apply Denes to Shoubridge is further weakened because Shoubridge does not introduce new edges in his simulations, while Denes requires the removal of known existing edges and adding new non-previously existing edges. Indeed, Dr. Karger agreed that “[Shoubridge] is only going to have those edges fail or come

back, and he's not going to introduce other edges in this simulation...." Ex. 2010, 11/29/16 Karger Tr. 51:20-23; *See* Denes at 366 (describing the disconnection of edges connected to  $p$  and adding edges between pairs of nodes that were previously connected to  $p$ ). Further, a POSITA would not have sought to add the teachings of Denes in Shoubridge because the simulation of a dynamic network was sufficient for its intended purpose, to test a specific hybrid routing algorithm, and would complicate Shoubridge's simulation. Moreover, Shoubridge's simulation of a dynamic network was based on simulating the links, not adding or removal of nodes.

146. Similarly, a POSITA would not have sought to add Rufino into the mix because there were no actual nodes to remove nor does Rufino relate to the overall simulation or the hybrid routing algorithm. Rather Rufino relates to a token bus network, where messages are *only* sent on a physical bus and the logical ring established in Rufino is at best a 2-regular network. Petitioner recognizes that "Rufino relates to token ring which is a 2-regular topology." Petition at 59; *see also id.* at 16 ("Rufino discloses messaging techniques for maintaining a 2-regular non-complete network when a node is disconnected from the network."). In other words, Rufino is directed towards 2-regular systems not "a broadcast channel, said broadcast channel forming an  $m$ -regular graph where  $m$  is at least 3" as required by Claim 6. Notably, and as discussed above, Petitioner's arguments center around

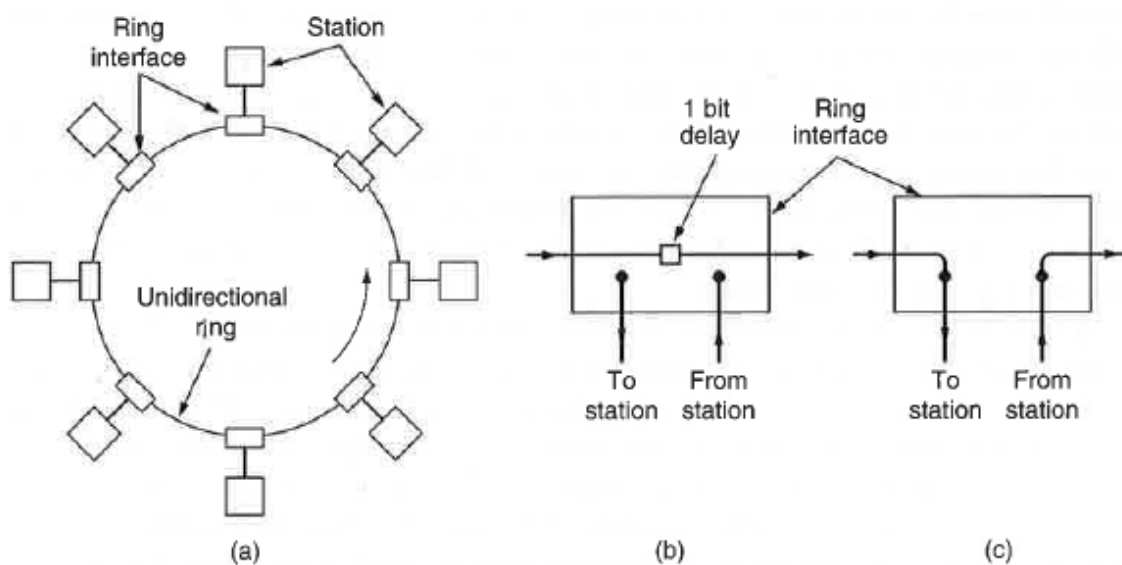
Rufino's discussion of "who\_follows" and "set\_successor" frames. *See* Petition at 32–37 (citing Rufino at 963 ("During this *who follows* query the current token holder waits for a *set\_successor* frame, until the end of a three-slot-time period.")). The "who\_follows" query involves a station recognizing that its successor station has left the logical token ring and broadcasting a message to all other stations connected to the physical bus that identifies the departed station. *See* Tanenbaum at 292. When the departed station's "sees a WHO\_FOLLOWS frame naming its predecessor, it responds by sending a SET\_SUCCESSOR frame to the station whose successor failed, naming itself as the new successor." *Id.* These operations work in the context of an  $m=2$  logical token ring precisely because each station knows its predecessor and successor stations, and there is only a single station defined as the successor of the failed station. This messaging procedure is not generalizable to  $m$ -regular networks in which  $m$  is at least 3 at least because once  $m$  exceeds 2 there are no longer "predecessor" and "successor" stations who can reply to a WHO\_FOLLOWS query.

147. As discussed above, and admitted by Dr. Karger, these frames are sent on a physical bus as opposed to a token ring. Petitioner does not provide an explanation of how Rufino's physical bus messaging system could be implemented to perform Denes'  $ET^{-1}$  transformation in an  $m$ -regular network in which  $m \geq 3$ . *See* Petition 33, 37-38. That is because Denes' transformation equation does not

work on a physical bus system, as each node is physically connected to each other (typically by copper wire). Moreover, Rufino, teaches away from the specialized point-to-point graphs solutions proposed by Shoubridge and Denes by stating because these solutions are “costly and complex.” *See* Rufino at 958 (criticizing specialized graph solutions such as “costly and complex.”) Thus, these three systems operate in fundamentally different ways that would completely redesign Rufino’s system to be contrary to its stated goals.

148. In addition, a POSITA would not have sought to add Todd in the mix of Shoubridge, Denes, and Rufino because Todd cannot use Rufino’s messaging system in Todd’s token grid network. Petitioner adds Todd in an attempt to cure the previously noted deficiencies of Rufino, namely how to extend Rufino “from a 2-regular topology to a 4-regular topology.” Petition at 58. But Todd is deficient. In particular, Todd is discussing token ring network as defined by IEEE 802.5 or FDDI while Rufino is discussing a token bus, as standardized in ISO 8802/4. Moreover, Todd’s 4-regular topology relies on modifying the physical connections of each station using an optical bypass switch to modify the pathways (i.e., switching from the DR state to the SR state, which requires no message passing at all) as opposed to the messaging system in Rufino and ISO 8802/4. *See* Todd at 8, Figs. 2, 3, and 4.

149. In a token ring network, the “ring is not really a broadcast medium, but a collection of individual point-to-point links that happen to form a circle.” Tanenbaum at 298. The topology of a token ring network, in which subsequent stations’ ring interfaces are connected by individual point-to-point connection is illustrated below:



**Fig. 4-28.** (a) A ring network. (b) Listen mode. (c) Transmit mode.

Tanenbaum at 293.

150. Unlike Rufino’s token bus network, token ring networks do not establish a logical ring over an underlying broadcast bus. *See* Tanenbaum at 287, 292. Rather, in a token ring network the “ring” refers to the physical connections between the stations themselves. Tanenbaum at 292. Todd’s “token grid network” only differs from the traditional token ring network defined in IEEE 802.5 or FDDI in that it is multidimensional. Todd at 1, col. 2 (“The token grid is a

multidimensional extension of the token ring [2], [4] where stations share access to a mesh formed by a set of overlapping rings.” (referencing [2] “ANSI/IEEE Standard 802.5, “Token ring access method and physical layer specifications, 1985”)). Thus, given the incompatibility of the topologies of Rufino’s token bus network and Todd’s token ring network, Petitioner is simply incorrect to say that “Todd... shows that the modification to Rufino can be done ‘in a very simple fashion’ by overlapping token rings.” Petition at 60. Indeed, Petitioner fails even to address the differences between token bus networks and token ring networks—misleadingly characterizing the Rufino reference as “relat[ing] to a token ring” when it is much more complex than what Petitioner seems to think. Petition at 59.

151. The incompatibility of Rufino’s token bus network and Todd’s modified token ring-to-grid network may be best understood with reference to Petitioner’s argument that Todd teaches how “to expand [Rufino’s] messaging techniques from a 2-regular topology to a 4-regular topology.” Petition at 59. Specifically, the way that Rufino and Todd handle the issue of failed nodes is completely different, with Todd’s technique not relying on messaging whatsoever. Rufino discloses using the `who_follows` and `set_successor` control frames in order to define the failed node’s successor as its predecessor’s new successor. In stark contrast, in the Todd token grid network, “optical bypass switches are used to maintain ring integrity following a station power failure.” Todd at 8, col. 2. That



is, Todd uses the optical bypass physical switches to simply bypass a failed station (in the DR or SR configuration) rather than using messages. In fact, Dr. Karger conceded that token bus network and token grid networks use different procedures for recovering from a failed station: ·

Q. · Is the procedure for recovering from a failed station different in token-bus networks and token ring networks?

A. · Well, there are a number of different recovery procedures that one could imagine. Rufino describes a particular recovery mechanism for the token-bus network. Todd describes another recovery mechanism which involves sort of shortcutting a failed station in order to recover from that failure.

Ex. 2010, Karger Tr., 75:9–18.

152. Accordingly, in my opinion POSITA would not be motivated to combine Shoubridge, Denes, Rufino, and Todd for the reasons stated above.

**DECLARATION**

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 by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed in Irvine, California on this 12th day of December, 2016.



Michael Goodrich, Ph.D.

# **APPENDIX A**

## CURRICULUM VITAE

Michael T. Goodrich

Dept. of Computer Science  
Bren School of Info. & Computer Sciences  
University of California, Irvine  
Irvine, CA 92697-3435

E-mail: goodrich (at) ieee.org  
<http://www.ics.uci.edu/~goodrich/>  
Phone: (949)824-9366  
Fax: (949)824-4056

**CITIZENSHIP:** U.S.A.

### EDUCATION

Ph.D.	1987	<i>Efficient Parallel Techniques for Computational Geometry</i> Computer Sciences, Purdue Univ. (M.J. Atallah, advisor)
M.S.	1985	Computer Sciences, Purdue Univ.
B.A.	1983	Mathematics and Computer Science, Calvin College

### PROFESSIONAL EXPERIENCE

April '07 to present	Chancellor's Professor, Dept. of Computer Science Univ. of California, Irvine
July '12 to June '13	Chair, Dept. of Computer Science Univ. of California, Irvine
October '06 to June '12	Assoc. Dean for Faculty Dev., Bren School of Info. and Comp. Sci. Univ. of California, Irvine
July '01 to March '07	Professor, Dept. of Computer Science Univ. of California, Irvine
Fall '00	Visiting Professor of Computer Science Brown Univ.
July '96 to June '02	Professor of Computer Science (on leave, from July '01) Johns Hopkins Univ.
July '92 to June '96	Associate Professor of Computer Science Johns Hopkins Univ.
Spring '94	Visiting Associate Professor of Computer Science Univ. of Illinois, Urbana-Champaign
July '87 to June '92	Assistant Professor of Computer Science Johns Hopkins Univ.

### RESEARCH INTERESTS

Algorithm and Data Structure Design  
Information Assurance and Security  
Parallel and Distributed Computing  
Graph and Geometric Algorithms

### HONORS AND AWARDS

- *Compere Loveless Fellowship in Computer Sciences*, Purdue Univ., 1985
- *Research Initiation Award*, National Science Foundation, 1988
- *Oraculum Award for Excellence in Teaching*, Johns Hopkins, 1993, 1994, 1995
- *ACM Recognition of Service Award*, 1996
- *Robert B. Pond, Sr. Award for Excellence in Undergraduate Teaching*, Johns Hopkins, 1998
- *Elected Senior Member*, the Institute of Electrical and Electronics Engineers (IEEE), 1999
- *Spirit of Technology Transition Award*, DARPA Dynamic Coalitions Program, 2002

- *Brown Univ. Award for Technological Innovation* (with R. Tamassia, N. Triandopoulos, D. Yao, and D. Ellis), 2006
- *ACM Distinguished Scientist*, 2006
- *2006 IEEE Computer Society Technical Achievement Award*, “for outstanding contributions to the design of parallel and distributed algorithms for fundamental combinatorial and geometric problems”
- *Fulbright Scholar*, 2007, for senior specialist service to University of Aarhus, Denmark
- *Fellow of the San Diego Supercomputer Center*, 2007
- *Fellow of the American Association for the Advancement of Science (AAAS)*, “for distinguished contributions to parallel and distributed algorithms for combinatorial and geometric problems, and excellence in teaching, academic and professional service, and textbook writing,” 2007
- *Fellow of the Institute of Electrical and Electronics Engineers (IEEE)*, “for contributions to parallel and distributed algorithms for combinatorial and geometric problems,” 2009
- *Fellow of the ACM*, “for contributions to data structures and algorithms for combinatorial and geometric problems,” 2009
- *ICS Dean’s Award for Research*, “for his contributions in the area of parallel and distributed algorithms,” 2014
- *Chancellor’s Award for Excellence in Fostering Undergraduate Research*, Univ. of California, Irvine, 2016
- *Faculty Mentor of the Month*, Undergraduate Research Opportunities Program (UROP), Univ. of California, Irvine, April 2016

## PUBLICATIONS

### Patents and Patent Applications:

- P-1. G. Ateniese, B. de Medeiros, and M.T. Goodrich, “Intermediated Delivery Scheme for Asymmetric Fair Exchange of Electronic Items,” U.S. Patent Application US 2004/0073790 A1, April 15, 2004.
- P-2. M.T. Goodrich and R. Tamassia, “Efficient Authenticated Dictionaries with Skip Lists and Commutative Hashing,” U.S. Patent 7,257,711, August 14, 2007.
- P-3. J.W. Green, J.L. Schultz, Y. Amir, and M.T. Goodrich, “High Refresh-Rate Retrieval of Freshly Published Content using Distributed Crawling,” U.S. Patent 7,299,219, November 20, 2007.
- P-4. R. Tamstorf, M.T. Goodrich, D. Eppstein, “Attribute Transfer Between Computer Models Including Identifying Isomorphic Regions in Polygonal Meshes,” U.S. Patent 8,681,145, March 25, 2014. (also Application US 2010/0238166 A1, September 23, 2010).
- P-5. N. Triandopoulos, M.T. Goodrich, D. Nguyen, O. Ohrimenko, C. Papamanthou, R. Tamassia, C.V. Lopes, “Techniques for Verifying Search Results Over a Distributed Collection,” U.S. Patent, 9,152,716, October 6, 2015.

### Books and Monographs:

- B-1. M.T. Goodrich and R. Tamassia, *Data Structures and Algorithms in Java*, John Wiley and Sons, Inc., 1998.
- B-2. M.T. Goodrich and C.C. McGeoch, eds., *Algorithm Engineering and Experimentation*, Lecture Notes in Computer Science (LNCS), Vol. 1619, Springer-Verlag, 1999.

- B-3. M.T. Goodrich and R. Tamassia, *Data Structures and Algorithms in Java, Second Edition*, John Wiley and Sons, Inc., 2001.
- B-4. M.T. Goodrich and R. Tamassia, *Algorithm Design: Foundations, Analysis, and Internet Examples*, John Wiley and Sons, Inc., 2002.
- B-5. M.T. Goodrich and S.G. Kobourov, eds., *10th Int. Symp. on Graph Drawing (GD)*, Lecture Notes in Computer Science, Vol. 2528, Springer-Verlag, 2002.
- B-6. M.T. Goodrich, R. Tamassia, and D. Mount, *Data Structures and Algorithms in C++*, John Wiley and Sons, Inc., 2004.
- B-7. M.T. Goodrich and R. Tamassia, *Data Structures and Algorithms in Java, Third Edition*, John Wiley and Sons, Inc., 2004.
- B-8. M.T. Goodrich and R. Tamassia, *Data Structures and Algorithms in Java, Fourth Edition*, John Wiley and Sons, Inc., 2006.
- B-9. M.T. Goodrich and R. Tamassia, *Data Structures and Algorithms in Java, Fifth Edition*, John Wiley and Sons, Inc., 2011.
- B-10. M.T. Goodrich and R. Tamassia, *Introduction to Computer Security*, Addison-Wesley, Inc., 2011.
- B-11. M.T. Goodrich, R. Tamassia, and D. Mount, *Data Structures and Algorithms in C++, Second Edition*, John Wiley and Sons, Inc., 2011.
- B-12. M.T. Goodrich, R. Tamassia, and M. Goldwasser, *Data Structures and Algorithms in Python*, John Wiley and Sons, Inc., 2013.
- B-13. M.T. Goodrich, R. Tamassia, and M. Goldwasser, *Data Structures and Algorithms in Java, Sixth Edition*, John Wiley and Sons, Inc., 2014.
- B-14. M.T. Goodrich and R. Tamassia, *Algorithm Design and Applications*, Wiley, 2015.

#### **Book Chapters:**

- Ch-1. M.J. Atallah and M.T. Goodrich, “Deterministic Parallel Computational Geometry,” in *Synthesis of Parallel Algorithms*, J.H. Reif, ed., Morgan Kaufmann, 497–536, 1993.
- Ch-2. M.T. Goodrich, “The Grand Challenges of Geometric Computing,” in *Developing a Computer Science Agenda for High-Performance Computing*, U. Vishkin, ed., ACM Press, 64–68, 1994.
- Ch-3. M.T. Goodrich, “Parallel Algorithms in Geometry,” *CRC Handbook of Discrete and Computational Geometry*, J.E. Goodman and J. O’Rourke, eds., CRC Press, Inc., 669–682, 1997.
- Ch-4. M.T. Goodrich and K. Ramaiyer, “Geometric Data Structures,” *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, eds., Elsevier Science Publishing, 463–489, 2000.
- Ch-5. M.T. Goodrich and R. Tamassia, “Simplified Analyses of Randomized Algorithms for Searching, Sorting, and Selection,” *Handbook of Randomized Computing*, S. Rajasekaran, P.M. Pardalos, J.H. Reif, and J.D.P. Rolim, eds., Kluwer Academic Publishers, Vol. 1, 23–34, 2001.
- Ch-6. M.T. Goodrich, “Parallel Algorithms in Geometry,” *Handbook of Discrete and Computational Geometry, Second Edition*, J.E. Goodman and J. O’Rourke, eds., Chapman & Hall/CRC Press, Inc., 953–967, 2004. (Revised version of Ch-3.)
- Ch-7. C. Duncan and M.T. Goodrich, “Approximate Geometric Query Structures,” *Handbook of Data Structures and Applications*, Chapman & Hall/CRC Press, Inc., 26-1–26-17, 2005.

- Ch-8. M.T. Goodrich, R. Tamassia, and L. Vismara, “Data Structures in JDSL,” *Handbook of Data Structures and Applications*, Chapman & Hall/CRC Press, Inc., 43-1–43-22, 2005.
- Ch-9. Y. Cho, L. Bao and M.T. Goodrich, “Secure Location-Based Access Control in WLAN Systems,” *From Problem Toward Solution: Wireless and Sensor Networks Security*, Zhen Jiang and Yi Pan, eds., Nova Science Publishers, Inc., Chapter 17, 2007.
- Ch-10. M.T. Goodrich and M.J. Nelson, “Distributed Peer-to-Peer Data Structures,” *Handbook of Parallel Computing: Models, Algorithms and Applications*, R. Rajasekaran and J. Reif, eds., CRC Press, 17-1–17-17, 2008.
- Ch-11. C.A. Duncan and M.T. Goodrich, “Planar Orthogonal and Polyline Drawing Algorithms,” *Handbook of Graph Drawing and Visualization*, CRC Press, Inc., 223–246, 2013.

#### Journal Papers:

- J-1. M.J. Atallah and M.T. Goodrich, “Efficient Parallel Solutions to Some Geometric Problems,” *Journal of Parallel and Distributed Computing*, **3**(4), 1986, 492–507.
- J-2. M.T. Goodrich, “Finding the Convex Hull of a Sorted Point Set in Parallel,” *Information Processing Letters*, **26**, 1987, 173–179.
- J-3. H. ElGindy and M.T. Goodrich, “Parallel Algorithms for Shortest Path Problems in Polygons,” *The Visual Computer*, **3**(6), 1988, 371–378.
- J-4. M.J. Atallah and M.T. Goodrich, “Parallel Algorithms For Some Functions of Two Convex Polygons,” *Algorithmica*, **3**, 1988, 535–548.
- J-5. M.J. Atallah, R. Cole, and M.T. Goodrich, “Cascading Divide-and-Conquer: A Technique for Designing Parallel Algorithms,” *SIAM Journal on Computing*, **18**(3), 1989, 499–532.
- J-6. M.T. Goodrich, “Triangulating a Polygon in Parallel,” *Journal of Algorithms*, **10**, 1989, 327–351.
- J-7. M.T. Goodrich and M.J. Atallah, “On Performing Robust Order Statistics in Tree-Structured Dictionary Machines,” *Journal of Parallel and Distributed Computing*, **9**(1), 1990, 69–76.
- J-8. M.T. Goodrich and J.S. Snoeyink, “Stabbing Parallel Segments with a Convex Polygon,” *Computer Vision, Graphics and Image Processing*, **49**, 1990, 152–170.
- J-9. J. Johnstone and M.T. Goodrich, “A Localized Method for Intersecting Plane Algebraic Curve Segments,” *The Visual Computer*, **7**(2–3), 1991, 60–71.
- J-10. M.T. Goodrich, “Intersecting Line Segments in Parallel with an Output-Sensitive Number of Processors,” *SIAM Journal on Computing*, **20**(4), 1991, 737–755.
- J-11. R. Cole and M.T. Goodrich, “Optimal Parallel Algorithms for Point-Set and Polygon Problems,” *Algorithmica*, **7**, 1992, 3–23.
- J-12. M.T. Goodrich, “A Polygonal Approach to Hidden-Line and Hidden-Surface Elimination,” *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, **54**(1), 1992, 1–12.
- J-13. M.T. Goodrich, S. Shauck, and S. Guha, “Parallel Methods for Visibility and Shortest Path Problems in Simple Polygons,” *Algorithmica*, **8**, 1992, 461–486, with addendum in *Algorithmica*, **9**, 1993, 515–516.
- J-14. M.T. Goodrich, C. Ó’Dúnlaing, and C. Yap “Constructing the Voronoi Diagram of a Set of Line Segments in Parallel,” *Algorithmica*, **9**, 1993, 128–141.
- J-15. M.T. Goodrich, “Constructing the Convex Hull of a Partially Sorted Set of Points,” *Computational Geometry: Theory and Applications*, **2**, 1993, 267–278.

- J-16. M.T. Goodrich, “Constructing Arrangements Optimally in Parallel,” *Discrete and Computational Geometry*, **9**, 1993, 371–385.
- J-17. M.T. Goodrich, M.J. Atallah, and M. Overmars, “Output-Sensitive Methods for Rectilinear Hidden Surface Removal,” *Information and Computation*, **107**(1), 1993, 1–24.
- J-18. M.J. Atallah, P. Callahan, and M.T. Goodrich, “P-Complete Geometric Problems,” *Int. Journal of Computational Geometry & Applications*, **3**(4), 1993, 443–462.
- J-19. M.J. Atallah, M.T. Goodrich, and S.R. Kosaraju, “Parallel Algorithms for Evaluating Sequences of Set-Manipulation Operations,” *Journal of the ACM*, **41**(6), 1994, 1049–1088.
- J-20. M.T. Goodrich, “Efficient Piecewise-Linear Function Approximation Using the Uniform Metric,” *Discrete and Computational Geometry*, **14**, 1995, 445–462.
- J-21. H. Brönnimann and M.T. Goodrich, “Almost Optimal Set Covers in Finite VC-Dimension,” *Discrete and Computational Geometry*, **14**, 1995, 463–479.
- J-22. M.T. Goodrich, “Planar Separators and Parallel Polygon Triangulation,” *J. Computer and System Sciences*, **51**(3), 1995, 374–389.
- J-23. M.T. Goodrich, M. Ghouse, and J. Bright, “Sweep Methods for Parallel Computational Geometry,” *Algorithmica*, **15**(2), 1996, 126–153.
- J-24. M.T. Goodrich and S.R. Kosaraju, “Sorting on a Parallel Pointer Machine with Applications to Set Expression Evaluation,” *Journal of the ACM*, **43**(2), 1996, 331–361.
- J-25. A. Garg, M.T. Goodrich, and R. Tamassia, “Planar Upward Tree Drawings with Optimal Area,” *International Journal of Computational Geometry & Applications*, **6**(3), 1996, 333–356.
- J-26. M.H. Nodine, M.T. Goodrich, and J.S. Vitter, “Blocking for External Graph Searching,” *Algorithmica*, **16**(2), 1996, 181–214.
- J-27. R. Cole, M.T. Goodrich, C. Ó Dúnlaing, “A Nearly Optimal Deterministic Parallel Voronoi Diagram Algorithm,” *Algorithmica*, **16**, 1996, 569–617.
- J-28. G. Das and M.T. Goodrich, “On the Complexity of Optimization Problems for 3-Dimensional Convex Polyhedra and Decision Trees,” *Computational Geometry: Theory and Applications*, **8**, 1997, 123–137.
- J-29. M.T. Goodrich and R. Tamassia, “Dynamic Ray Shooting and Shortest Paths in Planar Subdivisions via Balanced Geodesic Triangulations,” *J. Algorithms*, **23**, 1997, 51–73.
- J-30. M. Ghouse and M.T. Goodrich, “Fast Randomized Parallel Methods for Planar Convex Hull Construction,” *Computational Geometry: Theory and Applications*, **7**, 1997, 219–235.
- J-31. L.P. Chew, M.T. Goodrich, D.P. Huttenlocher, K. Kedem, J.M. Kleinberg, and D. Kravets, “Geometric Pattern Matching under Euclidean Motion,” *Computational Geometry: Theory and Applications*, **7**, 1997, 113–124.
- J-32. M.T. Goodrich and E.A. Ramos, “Bounded-Independence Derandomization of Geometric Partitioning with Applications to Parallel Fixed-Dimensional Linear Programming,” *Discrete & Computational Geometry*, **18**(4), 1997, 397–420.
- J-33. M.T. Goodrich, “An Improved Ray Shooting Method for Constructive Solid Geometry Models via Tree Contraction,” *International Journal of Computational Geometry & Applications*, **8**(1), 1998, 1–23.
- J-34. G. Barequet, A.J. Briggs, M.T. Dickerson, and M.T. Goodrich, “Offset-Polygon Annulus Placement Problems,” *Computational Geometry: Theory and Applications*, **11**(3–4), 1998–99, 125–141.



- J-35. M.T. Goodrich and R. Tamassia, “Dynamic Trees and Dynamic Point Location,” *SIAM J. Comput.*, **28**(2), 1999, 612–636.
- J-36. G. Barequet, S.S. Bridgeman, C.A. Duncan, M.T. Goodrich, and R. Tamassia, “GeomNet: Geometric Computing Over the Internet,” *IEEE Internet Computing*, **3**(2), 1999, 21–29.
- J-37. M.T. Goodrich, J.S.B. Mitchell, and M.W. Orletsky, “Approximate Geometric Pattern Matching Under Rigid Motion,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, **21**(4), 1999, 371–379.
- J-38. M.T. Goodrich, “Communication-Efficient Parallel Sorting,” *SIAM Journal on Computing*, **29**(2), 1999, 416–432.
- J-39. C.A. Duncan, M.T. Goodrich, S.G. Kobourov, “Balanced Aspect Ratio Trees and Their Use for Drawing Very Large Graphs,” *Journal of Graph Algorithms and Applications*, **4**(3), 2000, 19–46. Also available at [www.cs.brown.edu/publications/jgaa/](http://www.cs.brown.edu/publications/jgaa/).
- J-40. M.T. Goodrich and C.G. Wagner, “A Framework for Drawing Planar Graphs with Curves and Polylines,” *Journal of Algorithms*, **37**, 2000, 399–421.
- J-41. C.A. Duncan, M.T. Goodrich, S.G. Kobourov, “Balanced Aspect Ratio Trees: Combining the Benefits of  $k$ -D Trees and Octrees,” *J. Algorithms*, **38**, 2001, 303–333.
- J-42. G. Barequet, M. Dickerson, and M.T. Goodrich, “Voronoi Diagrams for Polygon-Offset Distance Functions,” *Discrete and Computational Geometry*, **25**(2), 2001, 271–291.
- J-43. C.C. Cheng, C.A. Duncan, M.T. Goodrich, and S.G. Kobourov, “Drawing Planar Graphs with Circular Arcs,” *Discrete and Computational Geometry*, **25**(3), 2001, 405–418.
- J-44. N.M. Amato, M.T. Goodrich, and E.A. Ramos, “A Randomized Algorithm for Triangulating a Simple Polygon in Linear Time,” *Discrete and Computational Geometry*, **26**(2), 2001, 245–265.
- J-45. R. Tamassia, M.T. Goodrich, L. Vismara, M. Handy, G. Shubina, R. Cohen, B. Hudson, R.S. Baker, N. Gelfand, and U. Brandes, “JDSL: The Data Structures Library in Java,” *Dr. Dobbs Journal*, **323**, 2001, 21–31.
- J-46. G. Barequet, D.Z. Chen, O. Daescu, M.T. Goodrich, and J.S. Snoeyink, “Efficiently Approximating Polygonal Paths in Three and Higher Dimensions,” *Algorithmica*, **33**(2), 2002, 150–167.
- J-47. T. Chan, M.T. Goodrich, S.R. Kosaraju, and R. Tamassia, “Optimizing Area and Aspect Ratio in Straight-Line Orthogonal Tree Drawings,” *Computational Geometry: Theory and Applications*, **23**(2), 2002, 153–162.
- J-48. C.A. Duncan, M.T. Goodrich, and S.G. Kobourov, “Planarity-Preserving Clustering and Embedding for Large Planar Graphs,” *Computational Geometry: Theory and Applications*, **24**(2), 2003, 95–114.
- J-49. A.L. Buchsbaum and M.T. Goodrich, “Three-Dimensional Layers of Maxima,” *Algorithmica*, **39**, 2004, 275–286.
- J-50. G. Barequet, M.T. Goodrich, and C. Riley, “Drawing Graphs with Large Vertices and Thick Edges,” *J. of Graph Algorithms and Applications (JGAA)*, **8**(1), 2004, 3–20.
- J-51. G. Barequet, M.T. Goodrich, A. Levi-Steiner, and D. Steiner, “Contour Interpolation by Straight Skeletons,” *Graphical Models (GM)*, **66**(4), 2004, 245–260.
- J-52. P. Gajer, M.T. Goodrich, and S.G. Kobourov, “A Multi-Dimensional Approach to Force-Directed Layouts of Large Graphs,” *Computational Geometry: Theory and Applications*, **29**(1), 3–18, 2004.

- J-53. G. Barequet, P. Bose, M.T. Dickerson, and M.T. Goodrich, “Optimizing a Constrained Convex Polygonal Annulus,” *J. of Discrete Algorithms (JDA)*, **3**(1), 1–26, 2005.
- J-54. A. Bagchi, A.L. Buchsbaum, and M.T. Goodrich, “Biased Skip Lists,” *Algorithmica*, **42**(1), 31–48, 2005.
- J-55. M. Dickerson, D. Eppstein, M.T. Goodrich, J. Meng, “Confluent Drawings: Visualizing Non-planar Diagrams in a Planar Way,” *J. of Graph Algorithms and Applications (JGAA)*, **9**(1), 31–52, 2005.
- J-56. A. Bagchi, A. Chaudhary, M.T. Goodrich, C. Li, and M. Shmueli-Scheuer, “Achieving Communication Efficiency through Push-Pull Partitioning of Semantic Spaces to Disseminate Dynamic Information,” *IEEE Trans. on Knowledge and Data Engineering (TKDE)*, **18**(10), 1352–1367, 2006.
- J-57. D. Eppstein, M.T. Goodrich, and J.Y. Meng, “Confluent Layered Drawings,” *Algorithmica*, **47**(4), 439–452, 2007.
- J-58. A. Bagchi, A. Chaudhary, D. Eppstein, and M.T. Goodrich, “Deterministic Sampling and Range Counting in Geometric Data Streams,” *ACM Transactions on Algorithms*, **3**(2), Article 16, 2007, 18 pages.
- J-59. D. Eppstein, M.T. Goodrich, and D. Hirschberg, “Improved Combinatorial Group Testing Algorithms for Real-World Problem Sizes,” *SIAM Journal on Computing*, **36**(5), 1360–1375, 2007.
- J-60. D. Eppstein, M.T. Goodrich, and J.Z. Sun, “Skip Quadrees: Dynamic Data Structures for Multidimensional Point Sets,” *Int. Journal on Computational Geometry and Applications*, **18**(1/2), 131–160, 2008.
- J-61. M.T. Goodrich, “Probabilistic Packet Marking for Large-Scale IP Traceback,” *IEEE/ACM Transactions on Networking*, **16**(1), 15–24, 2008.
- J-62. M.T. Goodrich and D.S. Hirschberg, “Improved Adaptive Group Testing Algorithms with Applications to Multiple Access Channels and Dead Sensor Diagnosis,” *Journal of Combinatorial Optimization*, **15**(1), 95–121, 2008.
- J-63. M.T. Goodrich, R. Tamassia, and D. Yao, “Notarized Federated ID Management and Authentication,” *Journal of Computer Security*, **16**(4), 399–418, 2008.
- J-64. M.T. Goodrich, “Pipelined Algorithms to Detect Cheating in Long-Term Grid Computations,” *Theoretical Computer Science*, **408**, 199–207, 2008.
- J-65. D. Eppstein, M.T. Goodrich, E. Kim, and R. Tamstorf, “Motorcycle Graphs: Canonical Quad Mesh Partitioning,” *Computer Graphics Forum*, special issue on papers from 6th European Symp. on Geometry Processing (SGP), **27**(6), 1477–1486, 2008.
- J-66. M.T. Goodrich, M. Sirivianos, J. Solis, C. Soriente, G. Tsudik, E. Uzun, “Using Audio in Secure Device Pairing,” *Int. J. Security and Networks*, **4**(1/2), 57–68, 2009.
- J-67. M.T. Goodrich, “On the Algorithmic Complexity of the Mastermind Game with Black-Peg Results,” *Information Processing Letters*, **109**, 675–678, 2009.
- J-68. D. Eppstein, M.T. Goodrich, E. Kim, and R. Tamstorf, “Approximate Topological Matching of Quad Meshes,” *The Visual Computer*, **25**(8), 771–783, 2009.
- J-69. D. Eppstein and M.T. Goodrich, “Succinct Greedy Geometric Routing Using Hyperbolic Geometry,” *IEEE Transactions on Computers*, **60**(11), 1571–1580, 2011. Posted online Dec. 2010, IEEE Computer Society Digital Library.
- J-70. D. Eppstein, M.T. Goodrich, and D. Strash, “Linear-Time Algorithms for Geometric Graphs

- with Sublinearly Many Edge Crossings,” *SIAM Journal on Computing*, **39**(8), 3814–3829, 2010.
- J-71. M.T. Goodrich, R. Tamassia, and N. Triandopoulos, “Efficient Authenticated Data Structures for Graph Connectivity and Geometric Search Problems,” *Algorithmica*, **60**(3), 505–552, 2011.
- J-72. D. Eppstein and M.T. Goodrich, “Straggler Identification in Round-Trip Data Streams via Newton’s Identities and Invertible Bloom Filters,” *IEEE Transactions on Knowledge and Data Engineering* (TKDE), **23**(2), 297–306, 2011.
- J-73. C.A. Duncan, M.T. Goodrich, S.G. Kobourov, “Planar Drawings of Higher-Genus Graphs,” *Journal of Graph Algorithms and Applications*, **15**(1), 7–32, 2011.
- J-74. M.T. Dickerson, M.T. Goodrich, T.D. Dickerson, and Y.D. Zhuo “Round-Trip Voronoi Diagrams and Doubling Density in Geographic Networks,” *Transactions on Computational Science*, M.L. Gavrilova et al. (Eds.), Vol. 14, LNCS 6970, 211–238, 2011.
- J-75. M.T. Goodrich, “Randomized Shellsort: A Simple Data-Oblivious Sorting Algorithm,” *Journal of the ACM*, **58**(6), Article No. 27, 2011.
- J-76. C.A. Duncan, D. Eppstein, M.T. Goodrich, S. Kobourov, and M. Nöllenburg, “Lombardi Drawings of Graphs,” *Journal of Graph Algorithms and Applications (JGAA)*, **16**(1), 85–108, 2012.
- J-77. E. Wolf-Chambers, D. Eppstein, M.T. Goodrich, and M. Löffler, “Drawing Graphs in the Plane with a Prescribed Outer Face and Polynomial Area,” *Journal of Graph Algorithms and Applications (JGAA)*, **16**(2), 243–259, 2012.
- J-78. M.T. Goodrich, D. Nguyen, O. Ohrimenko, C. Papamanthou, R. Tamassia, N. Triandopoulos, and C.V. Lopes, “Efficient Verification of Web-Content Searching Through Authenticated Web Crawlers,” *Proc. VLDB*, **5**(10):920-931, 2012.
- J-79. D. Eppstein, M.T. Goodrich, D. Strash, and L. Trott, “Extended Dynamic Subgraph Statistics Using h-Index Parameterized Data Structures,” *Theoretical Computer Science*, **447**, 44–52, 2012.
- J-80. M.T. Goodrich, “Learning Character Strings via Mastermind Queries, With a Case Study Involving mtDNA,” *IEEE Transactions on Information Theory*, **58**(11), 6726–6736, 2012.
- J-81. A.U. Asuncion and M.T. Goodrich, “Nonadaptive Mastermind Algorithms for String and Vector Databases, with Case Studies,” *IEEE Transactions on Knowledge and Data Engineering* (TKDE), **25**(1), 131–144, 2013.
- J-82. C.A. Duncan, D. Eppstein, M.T. Goodrich, S. Kobourov, and M. Nöllenburg, “Drawing Trees with Perfect Angular Resolution and Polynomial Area,” *Discrete & Computational Geometry*, **49**(2), 157–182, 2013.
- J-83. E. Angelino, M.T. Goodrich, M. Mitzenmacher and J. Thaler, “External Memory Multimaps,” *Algorithmica*, **67**(1), 23–48, 2013.
- J-84. D. Eppstein, M.T. Goodrich, M. Löffler, D. Strash and L. Trott, “Category-Based Routing in Social Networks: Membership Dimension and the Small-World Phenomenon,” *Theoretical Computer Science*, **514**, 96–104, 2013.
- J-85. M.T. Goodrich, “Spin-the-bottle Sort and Annealing Sort: Oblivious Sorting via Round-robin Random Comparisons,” *Algorithmica*, **68**(4), 835–858, 2014.
- J-86. Michael J. Bannister, William E. Devanny, David Eppstein, and M.T. Goodrich, “The Galois Complexity of Graph Drawing: Why Numerical Solutions are Ubiquitous for Force-Directed,

Spectral, and Circle Packing Drawings,” *Journal of Graph Algorithms and Applications*, **19**(2), 619–656, 2015.

# **Papers in Peer-Reviewed Proceedings:**

- C-1. M.J. Atallah and M.T. Goodrich, “Efficient Parallel Solutions to Geometric Problems,” *1985 IEEE Int. Conf. on Parallel Processing (ICPP)*, 411–417. (Preliminary version of J-1.)
- C-2. F. Berman, M.T. Goodrich, C. Koelbel, W. Robison, and K. Showell, “Prep-P: A Mapping Preprocessor for CHiP Computers,” *1985 IEEE Int. Conf. on Parallel Processing*, 731–733.
- C-3. M.J. Atallah and M.T. Goodrich, “Parallel Algorithms For Some Functions of Two Convex Polygons,” *24th Allerton Conf. on Communication, Control and Computing*, 1986, 758–767. (Preliminary version of J-4.)
- C-4. M.J. Atallah and M.T. Goodrich, “Efficient Plane Sweeping in Parallel,” *2nd ACM Symp. on Computational Geometry (SoCG)*, 1986, 216–225.
- C-5. M.T. Goodrich, “A Polygonal Approach to Hidden-Line Elimination,” *25th Allerton Conf. on Communication, Control, and Computing*, 1987, 849–858. (Preliminary version of J-12.)
- C-6. M.J. Atallah, R. Cole, and M.T. Goodrich, “Cascading Divide-and-Conquer: A Technique for Designing Parallel Algorithms,” *28th IEEE Symp. on Foundations of Computer Science (FOCS)*, 1987, 151-160. (Preliminary version of J-5.)
- C-7. M.J. Atallah, M.T. Goodrich, and S.R. Kosaraju, “Parallel Algorithms for Evaluating Sequences of Set-Manipulation Operations,” *3rd Aegean Workshop on Computing (AWOC), Lecture Notes in Computer Science (LNCS)*: 319, Springer-Verlag, 1988, 1–10. (Preliminary version of J-19.)
- C-8. R. Cole and M.T. Goodrich, “Optimal Parallel Algorithms for Polygon and Point-Set Problems,” *4th ACM Symp. on Computational Geometry (SoCG)*, 1988, 201–210. (Preliminary version of J-11.)
- C-9. M.T. Goodrich, “Intersecting Line Segments in Parallel with an Output-Sensitive Number of Processors,” *1989 ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, 127–137. (Preliminary version of J-10.)
- C-10. M.T. Goodrich and S.R. Kosaraju, “Sorting on a Parallel Pointer Machine with Applications to Set Expression Evaluation,” *30th IEEE Symp. on Foundations of Computer Science (FOCS)*, 1989, 190–195. (Preliminary version of J-24.)
- C-11. M.T. Goodrich, C. Ó’Dúnlaing, and C. Yap “Constructing the Voronoi Diagram of a Set of Line Segments in Parallel,” *Lecture Notes in Computer Science 382, Algorithms and Data Structures (WADS)*, Springer-Verlag, 1989, 12–23. (Preliminary version of J-14.)
- C-12. M.T. Goodrich and J.S. Snoeyink, “Stabbing Parallel Segments with a Convex Polygon,” *Lecture Notes in Computer Science 382, Algorithms and Data Structures (WADS)*, Springer-Verlag, 1989, 231–242. (Preliminary version of J-8.)
- C-13. J. Johnstone and M.T. Goodrich, “A Localized Method for Intersecting Plane Algebraic Curve Segments,” *New Advances in Computer Graphics: Proc. of Computer Graphics International ’89*, R.A. Earnshaw, B. Wyvel, eds., Springer-Verlag, 1989, 165–181. (Preliminary version of J-9.)
- C-14. M.J. Atallah, P. Callahan, and M.T. Goodrich, “P-Complete Geometric Problems,” *2nd ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, 1990, 317–326. (Preliminary version of J-18.)

- C-15. R. Cole, M.T. Goodrich, C. Ó Dúnlaing, “Merging Free Trees in Parallel for Efficient Voronoi Diagram Construction”, *17th Int. Conf. on Automata, Languages, and Programming* (ICALP), 1990, 432–445. (Preliminary version of J-27.)
- C-16. M.T. Goodrich, M.J. Atallah, and M. Overmars, “An Input-Size/Output-Size Trade-Off in the Time-Complexity of Rectilinear Hidden-Surface Removal”, *17th Int. Conf. on Automata, Languages, and Programming* (ICALP), 1990, 689–702. (Preliminary version of J-17.)
- C-17. M.T. Goodrich, M. Ghouse, and J. Bright, “Generalized Sweep Methods for Parallel Computational Geometry,” *2nd ACM Symp. on Parallel Algorithms and Architectures* (SPAA), 1990, 280–289. (Preliminary version of J-23.)
- C-18. M.T. Goodrich, “Applying Parallel Processing Techniques to Classification Problems in Constructive Solid Geometry,” *1st ACM-SIAM Symp. on Discrete Algorithms* (SODA), 1990, 118–128. (Preliminary version of J-33.)
- C-19. M.T. Goodrich, S. Shauck, and S. Guha, “Parallel Methods for Visibility and Shortest Path Problems in Simple Polygons,” *6th ACM Symp. on Computational Geometry* (SoCG), 1990, 73–82. (Preliminary version of J-13.)
- C-20. M. Ghouse and M.T. Goodrich, “In-Place Techniques for Parallel Convex Hull Algorithms,” *3rd ACM Symp. on Parallel Algorithms and Architectures* (SPAA), 1991, 192–203. (Preliminary version of J-30.)
- C-21. M.T. Goodrich, “Constructing Arrangements Optimally in Parallel,” *3rd ACM Symp. on Parallel Algorithms and Architectures* (SPAA), 1991, 169–179. (Preliminary version of J-16.)
- C-22. M.T. Goodrich and R. Tamassia, “Dynamic Trees and Dynamic Point Location,” *23rd ACM Symp. on Theory of Computing* (STOC), 1991, 523–533. (Preliminary version of J-35.)
- C-23. M.T. Goodrich, “Using Approximation Algorithms to Design Parallel Algorithms that May Ignore Processor Allocation,” *32nd IEEE Symp. on Foundations of Computer Science* (FOCS), 1991, 711–722.
- C-24. M.T. Goodrich, “Planar Separators and Parallel Polygon Triangulation,” *24th ACM Symp. on Theory of Computing* (STOC), 1992, 507–516. (Preliminary version of J-22.)
- C-25. M.T. Goodrich, Y. Matias, U. Vishkin, “Approximate Parallel Prefix Computation and Its Applications,” *7th IEEE Int. Parallel Processing Symp* (IPPS), 1993, 318–325.
- C-26. M. Ghouse and M.T. Goodrich, “Experimental Evidence for the Power of Random Sampling in Practical Parallel Algorithms,” *7th IEEE Int. Parallel Processing Symp* (IPPS), 1993, 549–556.
- C-27. L.P. Chew, M.T. Goodrich, D.P. Huttenlocher, K. Kedem, J.M. Kleinberg, and D. Kravets, “Geometric Pattern Matching under Euclidean Motion,” *5th Canadian Conf. on Computational Geometry* (CCCG), 1993, 151–156. (Preliminary version of J-31.)
- C-28. M.T. Goodrich, “Geometric Partitioning Made Easier, Even in Parallel,” *9th ACM Symp. on Computational Geometry* (SoCG), 1993, 73–82.
- C-29. M.T. Goodrich and R. Tamassia, “Dynamic Ray Shooting and Shortest Paths via Balanced Geodesic Triangulations,” *9th ACM Symp. on Computational Geometry* (SoCG), 1993, 318–327. (Preliminary version of J-29.)
- C-30. A. Garg, M.T. Goodrich, and R. Tamassia, “Area-Efficient Upward Tree Drawings,” *9th ACM Symp. on Computational Geometry* (SoCG), 1993, 359–368. (Preliminary version of J-25.)
- C-31. M.H. Nodine, M.T. Goodrich, and J.S. Vitter, “Blocking for External Graph Searching,” *12th ACM Symp. on Principles of Database Systems* (PODS), 1993, 222–232. (Preliminary

version of J-26.)

- C-32. E.M. Arkin, M.T. Goodrich, J.S.B. Mitchell, D. Mount, and S.S. Skiena, "Point Probe Decision Trees for Geometric Concept Classes," *Lecture Notes in Computer Science 709: Algorithms and Data Structures* (WADS), Springer-Verlag, 1993, 95–106.
- C-33. M.T. Goodrich, J.J. Tsay, D.E. Vengroff, and J.S. Vitter, "External-Memory Computational Geometry," *34th IEEE Symp. on Foundations of Computer Science* (FOCS), 1993, 714–723.
- C-34. M.T. Goodrich, Y. Matias, and U. Vishkin, "Optimal Parallel Approximation Algorithms for Prefix Sums and Integer Sorting," *5th ACM-SIAM Symp. on Discrete Algorithms* (SODA), 1994, 241–250.
- C-35. H. Brönnimann and M.T. Goodrich, "Almost Optimal Set Covers in Finite VC-Dimension," *10th ACM Symp. on Computational Geometry* (SoCG), 1994, 293–302. (Preliminary version of J-21.)
- C-36. M.T. Goodrich, "Efficient Piecewise-Linear Function Approximation Using the Uniform Metric," *10th ACM Symp. on Computational Geometry* (SoCG), 1994, 322–331. (Preliminary version of J-20.)
- C-37. M.J. Atallah, M.T. Goodrich, and K. Ramaiyer, "Biased Finger Trees and Three-Dimensional Layers of Maxima," *10th ACM Symp. on Computational Geometry* (SoCG), 1994, 150–159.
- C-38. M.T. Goodrich, J.S.B. Mitchell, and M.W. Orletsky, "Practical Methods for Approximate Geometric Pattern Matching under Rigid Motions," *10th ACM Symp. on Computational Geometry* (SoCG), 1994, 103–112. (Preliminary version of J-37.)
- C-39. N.M. Amato, M.T. Goodrich, E.A. Ramos, "Parallel Algorithms for Higher-Dimensional Convex Hulls," *35th IEEE Symp. on Foundations of Computer Science* (FOCS), 1994, 683–694.
- C-40. P.J. Tanenbaum, M.T. Goodrich, and E.R. Scheinerman, "Characterization and Recognition of Point-Halfspace and Related Orders," *2nd Int. Symp. on Graph Drawing* (GD), Lecture Notes in Computer Science 894, Springer-Verlag, 1994, 234–245.
- C-41. Y.J. Chiang, M.T. Goodrich, E.F. Grove, R. Tamassia, D.E. Vengroff, and J.S. Vitter, "External-Memory Graph Algorithms," *6th ACM-SIAM Symp. on Discrete Algorithms* (SODA), 1995, 139–149.
- C-42. N.M. Amato, M.T. Goodrich, and E.A. Ramos, "Computing Faces in Segment and Simplex Arrangements," *27th ACM Symp. on Theory of Computing* (STOC), 1995, 672–682.
- C-43. P. Callahan, M.T. Goodrich, and K. Ramaiyer, "Topology B-Trees and Their Applications," *1995 Workshop on Algorithms and Data Structures* (WADS), Lecture Notes in Computer Science 955, Springer-Verlag, 381–392.
- C-44. G. Das and M.T. Goodrich, "On the Complexity of Approximating and Illuminating Three-Dimensional Convex Polyhedra," *1995 Workshop on Algorithms and Data Structures* (WADS), Lecture Notes in Computer Science 955, Springer-Verlag, 74–85. (Preliminary version of J-28.)
- C-45. M.T. Goodrich, "Fixed-Dimensional Parallel Linear Programming via Relative  $\epsilon$ -Approximations," *7th ACM-SIAM Symp. on Discrete Algorithms* (SODA), 1996, 132–141. (Preliminary version of J-32.)
- C-46. M. Chrobak, M.T. Goodrich, and R. Tamassia, "Convex Drawings of Graphs in Two and Three Dimensions," *12th ACM Symp. on Computational Geometry* (SoCG), 1996, 319–328.
- C-47. M.T. Goodrich, "Communication-Efficient Parallel Sorting," *28th ACM Symp. on Theory of*

- Computing* (STOC), 1996, 247–256. (Preliminary version of J-38.)
- C-48. T. Chan, M.T. Goodrich, S.R. Kosaraju, and R. Tamassia, “Optimizing Area and Aspect Ratio in Straight-Line Orthogonal Tree Drawings,” *4th Int. Symp. on Graph Drawing* (GD), Lecture Notes in Computer Science 1190, Springer-Verlag, 1996, 63–75. (Preliminary version of J-47.)
  - C-49. M.T. Goodrich, “Randomized Fully-Scalable BSP Techniques for Multi-Searching and Convex Hull Construction,” *8th ACM-SIAM Symp. on Discrete Algorithms* (SODA), 1997, 767–776.
  - C-50. C.A. Duncan, M.T. Goodrich, and E.A. Ramos, “Efficient Approximation and Optimization Algorithms for Computational Metrology,” *8th ACM-SIAM Symp. on Discrete Algorithms* (SODA), 1997, 121–130.
  - C-51. M.T. Goodrich, M. Orletsky, and K. Ramaiyer, “Methods for Achieving Fast Query Times in Point Location Data Structures,” *8th ACM-SIAM Symp. on Discrete Algorithms* (SODA), 1997, 757–766.
  - C-52. M.T. Goodrich, L.J. Guibas, J. Hershberger, P.J. Tanenbaum, “Snap Rounding Line Segments Efficiently in Two and Three Dimensions,” *13th ACM Symp. on Computational Geometry* (SoCG), 1997, 284–293.
  - C-53. G. Barequet, S.S. Bridgeman, C.A. Duncan, M.T. Goodrich, and R. Tamassia, “Classical Computational Geometry in GeomNet,” *13th ACM Symp. on Computational Geometry* (SoCG), 1997, 412–414.
  - C-54. G. Barequet, A. Briggs, M. Dickerson, C. Dima, and M.T. Goodrich, “Animating the Polygon-Offset Distance Function,” *13th ACM Symp. on Computational Geometry* (SoCG), 1997, 479–480, and the *Video Review for the 13th ACM Symp. on Computational Geometry* (SoCG).
  - C-55. G. Barequet, A. Briggs, M. Dickerson, and M.T. Goodrich, “Offset-Polygon Annulus Placement Problems,” *1997 Workshop on Algorithms and Data Structures* (WADS), 1997, 378–391. (Preliminary version of J-34.)
  - C-56. G. Barequet, M. Dickerson, and M.T. Goodrich, “Voronoi Diagrams for Polygon-Offset Distance Functions,” *1997 Workshop on Algorithms and Data Structures* (WADS), 1997, 200–209. (Preliminary version of J-42.)
  - C-57. N. Gelfand, M.T. Goodrich, and R. Tamassia, “Teaching Data Structure Design Patterns,” *29th ACM SIGCSE Technical Symp. on Computer Science Education*, 1998, 331–335.
  - C-58. M.T. Goodrich and R. Tamassia, “Teaching the Analysis of Algorithms with Visual Proofs,” *29th ACM SIGCSE Technical Symp. on Computer Science Education*, 1998, 207–211.
  - C-59. G. Barequet, D.Z. Chen, O. Daescu, M.T. Goodrich, and J.S. Snoeyink, “Efficiently Approximating Polygonal Paths in Three and Higher Dimensions,” *1998 ACM Symp. on Computational Geometry* (SoCG), 1998, 317–326. (Preliminary version of J-46.)
  - C-60. M.T. Goodrich and C.G. Wagner, “A Framework for Drawing Planar Graphs with Curves and Polyines,” *6th Int. Symp. on Graph Drawing* (GD), Lecture Notes in Computer Science 1547, Springer-Verlag, 1998, 153–166. (Preliminary version of J-40.)
  - C-61. C.A. Duncan, M.T. Goodrich, S.G. Kobourov, “Balanced Aspect Ratio Trees and Their Use for Drawing Very Large Graphs,” *6th Int. Symp. on Graph Drawing* (GD), Lecture Notes in Computer Science 1547, Springer-Verlag, 1998, 111–124. (Preliminary version of J-39.)
  - C-62. M.T. Goodrich, M. Handy, B. Hudson, and R. Tamassia, “Abstracting Positional Information

- in Data Structures: Locators and Positions in JDSL,” *OOPSLA ’98 Technical Notes*, 1998.
- C-63. M.T. Goodrich and J.G. Kloss II, “Tiered Vector: An Efficient Dynamic Array for JDSL,” *OOPSLA ’98 Technical Notes*, 1998.
- C-64. M.T. Goodrich, M. Handy, B. Hudson, and R. Tamassia, “Accessing the Internal Organization of Data Structures in the JDSL Library,” *Int. Workshop on Algorithm Engineering and Experimentation (ALENEX)*, Springer-Verlag, Lecture Notes in Computer Science, Vol. 1619, 1999, 124–139.
- C-65. C.A. Duncan, M.T. Goodrich, S.G. Kobourov, “Balanced Aspect Ratio Trees: Combining the Benefits of  $k$ -D Trees and Octrees,” *10th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 1999, 300–309. (Preliminary version of J-41.)
- C-66. R.S. Baker, M. Boilen, M.T. Goodrich, R. Tamassia, and B.A. Stibel, “Testers and Visualizers for Teaching Data Structures,” *30th ACM SIGCSE Technical Symp. on Computer Science Education*, 1999, 261–265.
- C-67. M.T. Goodrich and R. Tamassia, “Using Randomization in the Teaching of Data Structures and Algorithms,” *30th ACM SIGCSE Technical Symp. on Computer Science Education*, 1999, 53–57. (Preliminary version of Ch-5.)
- C-68. G. Barequet, C. Duncan, M.T. Goodrich, S. Kumar, M. Pop, “Efficient Perspective-Accurate Silhouette Computation,” *15th ACM Symp. on Computational Geometry (SoCG)*, 1999, 417–418, and the *Video Review for the 15th ACM Symp. on Computational Geometry (SoCG)*.
- C-69. C.C. Cheng, C.A. Duncan, M.T. Goodrich, and S.G. Kobourov, “Drawing Planar Graphs with Circular Arcs,” *7th Int. Symp. on Graph Drawing (GD)*, Lecture Notes in Computer Science 1731, Springer-Verlag, 1999, 117–126. (Preliminary version of J-43.)
- C-70. C.A. Duncan, M.T. Goodrich, and S.G. Kobourov, “Planarity-Preserving Clustering and Embedding for Large Planar Graphs,” *7th Int. Symp. on Graph Drawing (GD)*, Lecture Notes in Computer Science 1731, Springer-Verlag, 1999, 186–196. (Preliminary version of J-48.)
- C-71. M.T. Goodrich and J.G. Kloss II, “Tiered Vectors: Efficient Dynamic Arrays for Rank-Based Sequences,” *1999 Workshop on Algorithms and Data Structures (WADS)*, Lecture Notes in Computer Science 1663, Springer-Verlag, 1999, 205–216.
- C-72. M.T. Goodrich, “Competitive Tree-Structured Dictionaries,” *11th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2000, 494–495.
- C-73. N.M. Amato, M.T. Goodrich, and E.A. Ramos, “Computing the Arrangement of Curve Segments: Divide-and-Conquer Algorithms via Sampling,” *11th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2000, 705–706.
- C-74. S. Bridgeman, M.T. Goodrich, S.G. Kobourov, and R. Tamassia, “PILOT: An Interactive Tool for Learning and Grading,” *31st ACM SIGCSE Technical Symp. on Computer Science Education*, 2000, 139–143.
- C-75. S. Bridgeman, M.T. Goodrich, S.G. Kobourov, and R. Tamassia, “SAIL: A System for Generating, Archiving, and Retrieving Specialized Assignments in LaTeX,” *31st ACM SIGCSE Technical Symp. on Computer Science Education*, 2000, 300–304.
- C-76. N.M. Amato, M.T. Goodrich, and E.A. Ramos, “Linear-Time Triangulation of a Simple Polygon Made Easier Via Randomization,” *16th ACM Symp. on Computational Geometry (SoCG)*, 2000, 201–212. (Preliminary version of J-44.)
- C-77. A.L. Buchsbaum, M.T. Goodrich, and J.R. Westbrook, “Range Searching Over Tree Cross Products,” *8th European Symp. on Algorithms (ESA)*, Lecture Notes in Computer



- Science 1879, Springer-Verlag, 2000, 120–131.
- C-78. C.A. Duncan, M.T. Dickerson, and M.T. Goodrich, “ $k$ -D Trees are Better When Cut on the Longest Side,” *8th European Symp. on Algorithms (ESA)*, Lecture Notes in Computer Science 1879, Springer-Verlag, 2000, 179–190.
  - C-79. P. Gajer, M.T. Goodrich, and S.G. Kobourov, “A Fast Multi-Dimensional Algorithm for Drawing Large Graphs,” *8th Int. Symp. on Graph Drawing (GD)*, Lecture Notes in Computer Science 1984, Springer-Verlag, 2001, 211–221. (Preliminary version of J-52.)
  - C-80. G. Ateniese, B. de Medeiros, and M.T. Goodrich, “TRICERT: A Distributed Certified E-mail Scheme,” *Network and Distributed Systems Security Symp. (NDSS)*, 2001, 47–56.
  - C-81. M.T. Goodrich and R. Tamassia, “Teaching Internet Algorithmics,” *32nd ACM SIGCSE Technical Symp. on Computer Science Education*, 2001, 129–133.
  - C-82. M. Pop, G. Barequet, C.A. Duncan, M.T. Goodrich, W. Hwang, and S. Kumar, “Efficient Perspective-Accurate Silhouette Computation and Applications,” *17th ACM Symp. on Computational Geometry (SoCG)*, 2001, 60–68.
  - C-83. M.T. Goodrich and R. Tamassia, “Implementation of an Authenticated Dictionary with Skip Lists and Commutative Hashing,” *DARPA Information Survivability Conf. & Exposition II (DISCEX)*, IEEE Press, 2001, 68–82.
  - C-84. A. Bagchi, A. Chaudhary, R. Garg, M.T. Goodrich, and V. Kumar, “Seller-Focused Algorithms for Online Auctioning,” *2001 Workshop on Algorithms and Data Structures (WADS)*, Lecture Notes in Computer Science 2125, Springer-Verlag, 2001, 135–147.
  - C-85. A. Anagnostopoulos, M.T. Goodrich, R. Tamassia, “Persistent Authenticated Dictionaries and Their Applications,” *Information Security Conf. (ISC)*, Lecture Notes in Computer Science 2200, Springer-Verlag, 2001, 379–393.
  - C-86. M.T. Goodrich, R. Tamassia, and J. Hasic, “An Efficient Dynamic and Distributed Cryptographic Accumulator,” *5th Information Security Conf. (ISC)*, Lecture Notes in Computer Science 2433, Springer-Verlag, 2002, 372–388.
  - C-87. A.L. Buchsbaum and M.T. Goodrich, “Three-Dimensional Layers of Maxima,” *10th European Symp. on Algorithms (ESA)*, Lecture Notes in Computer Science 2461, Springer-Verlag, 2002, 257–267. (Preliminary version of J-49.)
  - C-88. A. Bagchi, A.L. Buchsbaum, and M.T. Goodrich, “Biased Skip Lists,” *13th Int. Symp. on Algorithms and Computation (ISAAC)*, Lecture Notes in Computer Science 2518, Springer-Verlag, 2002, 1–13. (Preliminary version of J-54.)
  - C-89. M.T. Goodrich, “Efficient Packet Marking for Large-Scale IP Traceback,” *9th ACM Conf. on Computer and Communications Security (CCS)*, 2002, 117–126. (Preliminary version of J-61.)
  - C-90. M.T. Goodrich, R. Tamassia, N. Triandopoulos, and R. Cohen, “Authenticated Data Structures for Graph and Geometric Searching,” *RSA Conf.—Cryptographers’ Track (CT-RSA)*, Lecture Notes in Computer Science 2612, Springer-Verlag, 2003, 295–313. (Preliminary version of J-71.)
  - C-91. M.T. Goodrich, M. Shin, R. Tamassia, and W.H. Winsborough, “Authenticated Dictionaries for Fresh Attribute Credentials,” *1st Int. Conf. on Trust Management (iTrust)*, Lecture Notes in Computer Science 2692, Springer-Verlag, 2003, 332–347.
  - C-92. G. Barequet, M.T. Goodrich, A. Levi-Steiner, and D. Steiner, “Straight-Skeleton Based Contour Interpolation,” *14th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2003, 119–127. (Preliminary version of J-51.)

- C-93. G. Barequet, M.T. Goodrich, and C. Riley, "Drawing Graphs with Large Vertices and Thick Edges," *2003 Workshop and Data Structures and Algorithms (WADS)*, Lecture Notes in Computer Science 2748, Springer-Verlag, 2003, 281–293. (Preliminary version of J-50.)
- C-94. A. Bagchi, A. Chaudhary, M.T. Goodrich, and S. Xu, "Constructing Disjoint Paths for Secure Communication," *17th Int. Symp. on Distributed Computing (DISC)*, Lecture Notes in Computer Science 2848, Springer-Verlag, 2003, 181–195.
- C-95. M. Dickerson, D. Eppstein, M.T. Goodrich, J. Meng, "Confluent Drawings: Visualizing Non-planar Diagrams in a Planar Way," *11th Int. Symp. on Graph Drawing (GD)*, Lecture Notes in Computer Science 2912, Springer-Verlag, 2003, 1–12. (Preliminary version of J-55.)
- C-96. F. Brandenburg, D. Eppstein, M.T. Goodrich, S. Kobourov, G. Liotta, P. Mutzel, "Selected Open Problems in Graph Drawing," *11th Int. Symp. on Graph Drawing (GD)*, Lecture Notes in Computer Science 2912, Springer-Verlag, 2003, 515–539.
- C-97. A. Bagchi, A. Chaudhary, D. Eppstein, and M.T. Goodrich, "Deterministic Sampling and Range Counting in Geometric Data Streams," *20th ACM Symp. on Computational Geometry (SoCG)*, 144–151, 2004. (Preliminary version of J-58.)
- C-98. M.T. Goodrich, J.Z. Sun, and R. Tamassia, "Efficient Tree-Based Revocation in Groups of Low-State Devices," *Advances in Cryptology (CRYPTO)*, Springer, Lecture Notes in Computer Science 3152, 511–527, 2004.
- C-99. D. Eppstein, M.T. Goodrich, and J.Y. Meng, "Confluent Layered Drawings," *12th Int. Symp. on Graph Drawing (GD)*, Springer, Lecture Notes in Computer Science 3383, 184–194, 2004. (Preliminary version of J-57.)
- C-100. M.J. Atallah, K.B. Frikken, M.T. Goodrich, and R. Tamassia, "Secure Biometric Authentication for Weak Computational Devices," *9th Int. Conf. on Financial Cryptography and Data Security*, Springer, Lecture Notes in Computer Science 3570, 357–371, 2005.
- C-101. M.T. Goodrich, "Leap-Frog Packet Linking and Diverse Key Distributions for Improved Integrity in Network Broadcasts," *IEEE Symp. on Security and Privacy (S&P)*, 196–207, 2005.
- C-102. D. Eppstein, M.T. Goodrich, and J.Z. Sun, "The Skip Quadtree: A Simple Dynamic Data Structure for Multidimensional Data," *21st ACM Symp. on Computational Geometry (SoCG)*, 296–305, 2005. (Preliminary version of J-60.)
- C-103. M.J. Atallah, M.T. Goodrich, and R. Tamassia, "Indexing Information for Data Forensics," *3rd Applied Cryptography and Network Security Conf. (ACNS)*, Lecture Notes in Computer Science 3531, Springer, 206–221, 2005.
- C-104. W. Du and M.T. Goodrich, "Searching for High-Value Rare Events with Uncheatable Grid Computing," *3rd Applied Cryptography and Network Security Conf. (ACNS)*, Lecture Notes in Computer Science 3531, Springer, 122–137, 2005.
- C-105. L. Arge, D. Eppstein, and M.T. Goodrich, "Skip-Webs: Efficient Distributed Data Structures for Multi-Dimensional Data Sets," *24th ACM Symp. on Principles of Distributed Computing (PODC)*, 2005.
- C-106. D. Eppstein, M.T. Goodrich, and D. Hirschberg, "Improved Combinatorial Group Testing for Real-World Problem Sizes," *Workshop on Algorithms and Data Structures (WADS)*, Lecture Notes in Computer Science 3608, Springer, 86–98, 2005. (Preliminary version of J-59.)
- C-107. A. Chaudhary and M.T. Goodrich, "Balanced Aspect Ratio Trees Revisited," *Workshop on Algorithms and Data Structures (WADS)*, Lecture Notes in Computer Science 3608, Springer,

73–85, 2005.

- C-108. M.T. Goodrich, R. Tamassia, and D. Yao, “Accredited DomainKeys: A Service Architecture for Improved Email Validation,” *2nd Conf. on Email and Anti-Spam* (CEAS), 1–8, 2005.
- C-109. M.T. Goodrich, G.S. Lueker, and J.Z. Sun, “C-Planarity of Extrovert Clustered Graphs,” *13th Int. Symp. Graph Drawing* (GD), 211–222, 2005.
- C-110. D. Eppstein, M.T. Goodrich, J.Y. Meng, “Delta-Confluent Drawings,” *13th Int. Symp. Graph Drawing* (GD), 165–176, 2005.
- C-111. M.T. Goodrich, M.J. Nelson, and J.Z. Sun, “The Rainbow Skip Graph: A Fault-Tolerant Constant-Degree Distributed Data Structure,” *17th ACM-SIAM Symp. on Discrete Algorithms* (SODA), 384–393, 2006.
- C-112. M.T. Goodrich, M. Sirivianos, J. Solis, G. Tsudik, E. Uzun, “Loud And Clear: Human-Verifiable Authentication Based on Audio,” *26th IEEE Int. Conf. on Distributed Computing Systems* (ICDCS), 1–8, 2006. (Preliminary version of J-66.)
- C-113. M.T. Goodrich, R. Tamassia, and D. Yao, “Notarized Federated Identity Management for Web Services,” *20th IFIP WG Working Conf. on Data and Application Security* (DBSec), Springer, Lecture Notes in Computer Science, Vol. 4127, 133–147, 2006. (Preliminary version of J-63.)
- C-114. M.T. Goodrich and D.S. Hirschberg, “Efficient Parallel Algorithms for Dead Sensor Diagnosis and Multiple Access Channels,” *18th ACM Symp. on Parallelism in Algorithms and Architectures* (SPAA), 118–127, 2006. (Preliminary version of J-62.)
- C-115. Y. Cho, L. Bao, and M.T. Goodrich, “LAAC: A Location-Aware Access Control Protocol,” *Int. Workshop on Ubiquitous Access Control* (IWUAC), 1–7, 2006.
- C-116. M.B. Dillencourt, D. Eppstein, and M.T. Goodrich, “Choosing Colors for Geometric Graphs via Color Space Embeddings,” *14th Int. Symp. Graph Drawing* (GD), Lecture Notes in Computer Science, Vol. 4372, Springer, 294–305, 2006.
- C-117. D. Eppstein, M.T. Goodrich, and N. Sitchinava, “Guard Placement for Wireless Localization,” *23rd ACM Symp. on Computational Geometry* (SoCG), 27–36, 2007.
- C-118. M.T. Goodrich, C. Papamanthou, and R. Tamassia, “On the Cost of Persistence and Authentication in Skip Lists,” *6th Workshop on Experimental Algorithms* (WEA), LNCS 4525, 94–107, 2007.
- C-119. M.J. Atallah, M. Blanton, M.T. Goodrich, and S. Polu, “Discrepancy-Sensitive Dynamic Fractional Cascading, Dominated Maxima Searching, and 2-d Nearest Neighbors in Any Minkowski Metric,” *Workshop on Algorithms and Data Structures* (WADS), LNCS, Vol. 4619, Springer, 114–126, 2007.
- C-120. D. Eppstein and M.T. Goodrich, “Space-Efficient Straggler Identification in Round-Trip Data Streams via Newton’s Identities and Invertible Bloom Filters,” *Workshop on Algorithms and Data Structures* (WADS), LNCS, Vol. 4619, Springer, 2007, 638–649. (Preliminary version of J-72.)
- C-121. M.T. Goodrich and J.Z. Sun, “Checking Value-Sensitive Data Structures in Sublinear Space,” *18th Int. Symp. on Algorithms and Computation* (ISAAC), LNCS, vol. 4835, Springer, 2007, 353–364.
- C-122. M.T. Goodrich, R. Tamassia, and N. Triandopoulos, “Super-Efficient Verification of Dynamic Outsourced Databases,” *RSA Conf.—Cryptographers’ Track* (CT-RSA), LNCS, vol. 4964, Springer, 2008, 407–424.

- C-123. D. Eppstein, M.T. Goodrich, E. Kim, and R. Tamstorf, “Approximate Topological Matching of Quadrilateral Meshes,” *IEEE Int. Conf. on Shape Modeling and Applications (SMI)*, 2008, 83–92. (Preliminary version of J-68.)
- C-124. G. Barequet, D. Eppstein, M.T. Goodrich, and A. Waxman, “Straight Skeletons of Three-Dimensional Polyhedra,” *16th European Symp. on Algorithms (ESA)*, LNCS, vol. 5193, 2008, 148–160.
- C-125. M.T. Goodrich, C. Papamanthou, R. Tamassia, and N. Triandopoulos, “Athos: Efficient Authentication of Outsourced File Systems,” *11th Information Security Conf. (ISC)*, LNCS, vol. 5222, 2008, 80–96.
- C-126. L. Arge, M.T. Goodrich, M. Nelson, and N. Sitchinava, “Fundamental Parallel Algorithms for Private-Cache Chip Multiprocessors,” *20th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, 2008, 197–206.
- C-127. D. Eppstein and M.T. Goodrich, “Succinct Greedy Graph Drawing in the Hyperbolic Plane,” *16th Int. Symp. on Graph Drawing (GD)*, LNCS, vol. 5417, Springer, 2008, 14–25. (Preliminary version of J-69.)
- C-128. D. Eppstein and M.T. Goodrich, “Studying (Non-Planar) Road Networks Through an Algorithmic Lens,” *16th ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (GIS)*, 2008, 125–134. **Best Paper Award.**
- C-129. M. Dickerson and M.T. Goodrich, “Two-Site Voronoi Diagrams in Geographic Networks,” *16th ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (GIS)*, 2008, 439–442.
- C-130. D. Eppstein, M.T. Goodrich, and D. Strash, “Linear-Time Algorithms for Geometric Graphs with Sublinearly Many Crossings,” *20th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2009, 150–159. (Preliminary version of J-70.)
- C-131. M.T. Goodrich, “The Mastermind Attack on Genomic Data,” *30th IEEE Symp. on Security and Privacy (SSP)*, 2009, 204–218. (Preliminary version of J-80.)
- C-132. W. Du, D. Eppstein, M.T. Goodrich, and G.S. Lueker, “On the Approximability of Geometric and Geographic Generalization and the Min-Max Bin Covering Problem,” *Algorithms and Data Structures Symp. (WADS)*, LNCS, vol. 5664, Springer, 2009, 242–253.
- C-133. M.T. Goodrich, R. Tamassia, and N. Triandopoulos, J.Z. Sun, “Reliable Resource Searching in P2P Networks,” *5th Int. ICST Conf. on Security and Privacy in Communication Networks (SecureComm)*, Lecture Notes of ICST, vol. 19, Springer, 2009, 437–447.
- C-134. C.A. Duncan, M.T. Goodrich, S.G. Kobourov, “Planar Drawings of Higher-Genus Graphs,” *17th Int. Symp. on Graph Drawing (GD)*, LNCS, Springer, vol. 5849, 2009, 45–56. (Preliminary version of J-73.)
- C-135. D. Eppstein, M.T. Goodrich, L. Trott, “Going Off-road: Transversal Complexity in Road Networks,” *17th ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (GIS)*, 2009, 23–32.
- C-136. M.T. Goodrich and Darren Strash, “Succinct Greedy Geometric Routing in the Euclidean Plane,” *20th Int. Symp. on Algorithms and Computation (ISAAC)*, LNCS, vol. 5878, Springer, 2009, 781–791.
- C-137. M.T. Goodrich, “Randomized Shellsort: A Simple Oblivious Sorting Algorithm,” *21st ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2010, 1262–1277. (Preliminary version of J-75.)

- C-138. L. Arge, M.T. Goodrich, and N. Sitchinava, “Parallel External Memory Graph Algorithms,” *24th IEEE Int. Parallel & Distributed Processing Symp. (IPDPS)*, 2010, 1–11.
- C-139. G. Wang, T. Luo, M.T. Goodrich, W. Du, and Z. Zhu, “Bureaucratic Protocols for Secure Two-Party Sorting, Selection, and Permuting,” *5th ACM Symp. on Information, Computer and Communications Security*, 2010, 226–237.
- C-140. M.T. Dickerson, M.T. Goodrich, and T.D. Dickerson, “Round-Trip Voronoi Diagrams and Doubling Density in Geographic Networks,” *7th Int. Symp. on Voronoi Diagrams in Science and Engineering (ISVD)*, IEEE Press, 132–141, 2010. (Preliminary version of J-74.)
- C-141. M.T. Dickerson, D. Eppstein, and M.T. Goodrich, “Cloning Voronoi Diagrams via Retroactive Data Structures,” *18th European Symp. on Algorithms (ESA)*, LNCS, vol. 6346, 2010, 362–373.
- C-142. C.A. Duncan, D. Eppstein, M.T. Goodrich, S. Kobourov, and M. Nöllenburg, “Lombardi Drawings of Graphs,” *18th Int. Symp. on Graph Drawing (GD)*, LNCS, vol. 6502, 2010, 195–207. (Preliminary version of J-76.)
- C-143. E. Wolf-Chambers, D. Eppstein, M.T. Goodrich, and M. Löffler, “Drawing Graphs in the Plane with a Prescribed Outer Face and Polynomial Area,” *18th Int. Symp. on Graph Drawing (GD)*, LNCS, vol. 6502, 2010, 129–140. (Preliminary version of J-77.)
- C-144. C.A. Duncan, D. Eppstein, M.T. Goodrich, S. Kobourov, and M. Nöllenburg, “Drawing Trees with Perfect Angular Resolution and Polynomial Area,” *18th Int. Symp. on Graph Drawing (GD)*, LNCS, vol. 6502, 2010, 183–194. (Preliminary version of J-82.)
- C-145. A.U. Asuncion and M.T. Goodrich, “Turning Privacy Leaks into Floods: Surreptitious Discovery of Social Network Friendships and Other Sensitive Binary Attribute Vectors,” *Workshop on Privacy in the Electronic Society (WPES)*, held in conjunction with the 17th ACM Conf. on Computer and Communications Security (CCS), 2010, 21–30. (Preliminary version of J-81.)
- C-146. D. Eppstein, M.T. Goodrich, D. Strash, and L. Trott, “Extended Dynamic Subgraph Statistics Using h-Index Parameterized Data Structures,” *4th Annual Int. Conf. on Combinatorial Optimization and Applications (COCOAA)*, LNCS, vol. 6508, 2010, 128–141. (Preliminary version of J-79.)
- C-147. M.T. Goodrich and D. Strash, “Priority Range Trees,” *21st Int. Symp. on Algorithms and Computation (ISAAC)*, LNCS, vol. 6506, 2010, 97–108.
- C-148. D. Eppstein, M.T. Goodrich, R. Tamassia, “Privacy-Preserving Data-Oblivious Geometric Algorithms for Geographic Data,” *18th ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (GIS)*, 2010, 13–22.
- C-149. M.T. Goodrich, “Spin-the-bottle Sort and Annealing Sort: Oblivious Sorting via Round-robin Random Comparisons,” *8th Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, in conjunction with the ACM-SIAM Symp. on Discrete Algorithms (SODA), 2011. (Preliminary version of J-85.)
- C-150. M.T. Goodrich and F. Kerschbaum, “Privacy-Enhanced Reputation-Feedback Methods to Reduce Feedback Extortion in Online Auctions,” *ACM Conf. on Data and Application Security and Privacy (CODASPY)*, 2011, 273–282.
- C-151. M.T. Goodrich, “Data-Oblivious External-Memory Algorithms for the Compaction, Selection, and Sorting of Outsourced Data,” *23rd ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, 2011, 379–388.

- C-152. M.T. Goodrich and M. Mitzenmacher, “Brief Announcement: Large-Scale Multimaps,” *23rd ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, 2011, 259–260.
- C-153. D. Eppstein, M.T. Goodrich, F. Uyeda, and G. Varghese, “What’s the Difference? Efficient Set Synchronization without Prior Context,” *SIGCOMM* 218–229, 2011.
- C-154. D. Eppstein, M.T. Goodrich, and M. Löffler, “Tracking Moving Objects with Few Handovers,” *Algorithms and Data Structures Symp. (WADS)*, 362–373, LNCS, vol. 6844, 2011.
- C-155. M.T. Goodrich and M. Mitzenmacher, “Privacy-Preserving Access of Outsourced Data via Oblivious RAM Simulation,” *38th Int. Colloquium on Automata, Languages and Programming (ICALP)*, LNCS, vol. 6756, 2011, 576–587.
- C-156. M.T. Goodrich and P. Pszona, “External-Memory Network Analysis Algorithms for Naturally Sparse Graphs,” *European Symp. on Algorithms (ESA)*, LNCS, vol. 6942, 664–676, 2011.
- C-157. C. Duncan, D. Eppstein, M.T. Goodrich, S.G. Kobourov and M. Löffler, “Planar and Poly-Arc Lombardi Drawings,” *Int. Symp. Graph Drawing (GD)*, LNCS, vol. 7034, 308–319, 2011.
- C-158. R. Chernobelskiy, K. Cunningham, M.T. Goodrich, S.G. Kobourov and L. Trott, “Force-Directed Lombardi-Style Graph Drawing,” *Int. Symp. Graph Drawing (GD)*, LNCS, vol. 7034, 320–331, 2011.
- C-159. M.T. Goodrich and M. Mitzenmacher, “Invertible Bloom Lookup Tables,” *49th Allerton Conf. on Communication, Control, and Computing*, IEEE Press, **invited paper**, 2011.
- C-160. M.T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia, “Oblivious RAM Simulation with Efficient Worst-Case Access Overhead,” *ACM Cloud Computing Security Workshop (CCSW)*, in conjunction with the 17th ACM Conf. on Computer and Communications Security (CCS), 95–100, 2011.
- C-161. M.T. Goodrich and J.A. Simons, “Fully Retroactive Approximate Range and Nearest Neighbor Searching,” *22nd Int. Symp. on Algorithms and Computation (ISAAC)*, Springer, LNCS, vol. 7074, 292–301, 2011.
- C-162. E. Angelino, M.T. Goodrich, M. Mitzenmacher and J. Thaler, “External Memory Multimaps,” *22nd Int. Symp. on Algorithms and Computation (ISAAC)*, Springer, LNCS, vol. 7074, 384–394, 2011. (Preliminary version of J-83.)
- C-163. M.T. Goodrich, N. Sitchinava, and Q. Zhang, “Sorting, Searching, and Simulation in the MapReduce Framework,” *22nd Int. Symp. on Algorithms and Computation (ISAAC)*, Springer, LNCS, vol. 7074, 374–383, 2011.
- C-164. D. Eppstein, M.T. Goodrich, M. Löffler, D. Strash and L. Trott, “Category-Based Routing in Social Networks: Membership Dimension and the Small-World Phenomenon,” *IEEE Int. Conf. on Computational Aspects of Social Networks (CASoN)*, 102–107, 2011. (Preliminary version of J-84.)
- C-165. M.T. Goodrich, O. Ohrimenko, M. Mitzenmacher, and R. Tamassia, “Privacy-Preserving Group Data Access via Stateless Oblivious RAM Simulation,” *23rd ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 157–167, 2012.
- C-166. M.T. Goodrich, O. Ohrimenko, M. Mitzenmacher, and R. Tamassia, “Practical Oblivious Storage,” *2nd ACM Conf. on Data and Application Security and Privacy (CODASPY)*, 13–24, 2012.
- C-167. M.T. Goodrich and M. Mitzenmacher, “Anonymous Card Shuffling and its Applications

- to Parallel Mixnets,” *39th Int. Colloquium on Automata, Languages and Programming (ICALP)*, Springer, LNCS, vol. 6756, 576–587, 2012.
- C-168. M.T. Goodrich, O. Ohrimenko, and R. Tamassia, “Graph Drawing in the Cloud: Privately Visualizing Relational Data using Small Working Storage,” *20th Int. Symp. on Graph Drawing (GD)*, Springer, LNCS, vol. 7704, 43–54, 2012.
- C-169. F.J. Brandenburg, D. Eppstein, A. Gleissner, M.T. Goodrich, K. Hanauer, and J. Reislhuber, “On the Density of Maximal 1-Planar Graphs,” *20th Int. Symp. on Graph Drawing (GD)*, Springer, LNCS, vol. 7704, 327–338, 2012.
- C-170. M.J. Bannister, D. Eppstein, M.T. Goodrich, and L. Trott, “Force-Directed Graph Drawing Using Social Gravity and Scaling,” *20th Int. Symp. on Graph Drawing (GD)*, Springer, LNCS, vol. 7704, 414–425, 2012.
- C-171. M.T. Goodrich and J.A. Simons, “More Graph Drawing in the Cloud: Data-Oblivious st-Numbering, Visibility Representations, and Orthogonal Drawing of Biconnected Planar Graphs,” *20th Int. Symp. on Graph Drawing (GD)*, Springer, LNCS, vol. 7704, 569–570, 2012.
- C-172. M.T. Goodrich, D.S. Hirschberg, M. Mitzenmacher, and J. Thaler, “Cache-Oblivious Dictionaries and Multimaps with Negligible Failure Probability,” *Mediterranean Conf. on Algorithms (MedAlg)*, Springer, LNCS, vol. 7659, 203–218, 2012.
- C-173. D. Eppstein, M.T. Goodrich, and D.S. Hirschberg, “Combinatorial Pair Testing: Distinguishing Workers from Slackers,” *Algorithms and Data Structures Symp. (WADS)*, Springer, LNCS, vol. 8037, 316–327, 2013.
- C-174. D. Eppstein, M.T. Goodrich, and J.A. Simons, “Set-Difference Range Queries,” *25th Canadian Conf. on Computational Geometry (CCCG)*, 2013, <http://www.cccg.ca/proceedings/2013/>.
- C-175. M.T. Goodrich and P. Pszona, “Cole’s Parametric Search Technique Made Practical,” *25th Canadian Conf. on Computational Geometry (CCCG)*, 2013, <http://www.cccg.ca/proceedings/2013/>.
- C-176. L. Arge, M.T. Goodrich, F. van Walderveen, “Computing Betweenness Centrality in External Memory,” *IEEE Int. Conf. on Big Data (BigData)*, 368–375, 2013.
- C-177. M.T. Goodrich and P. Pszona, “Achieving Good Angular Resolution in 3D Arc Diagrams,” *21st Int. Symp. Graph Drawing (GD)*, Springer, LNCS, vol. 8242, 161–172, 2013.
- C-178. M.T. Goodrich and P. Pszona, “Streamed Graph Drawing and the File Maintenance Problem,” *21st Int. Symp. Graph Drawing (GD)*, Springer, LNCS, vol. 8242, 256–267, 2013.
- C-179. M.T. Goodrich, “Zig-zag Sort: A Simple Deterministic Data-Oblivious Sorting Algorithm Running in  $O(n \log n)$  Time,” *46th ACM Symp. on Theory of Computing (STOC)*, 684–693, 2014.
- C-180. D. Eppstein, M.T. Goodrich, M. Mitzenmacher, and P. Pszona, “Wear Minimization for Cuckoo Hashing: How Not to Throw a Lot of Eggs into One Basket,” *Symp. on Experimental Algorithms (SEA)*, Springer, LNCS, vol. 8504, 162–173, 2014.
- C-181. O. Ohrimenko, M.T. Goodrich, and R. Tamassia, and E. Upfal, “The Melbourne Shuffle: Improving Oblivious Storage in the Cloud,” *41st Int. Colloq. on Automata, Languages, and Programming (ICALP)*, Springer, LNCS, vol. 8573, 556–567, 2014.
- C-182. M.J. Bannister, W.E. Devanny, M.T. Goodrich, J.A. Simons, and Lowell Trott, “Windows into Geometric Events: Data Structures for Time-Windowed Querying of Temporal Point

- Sets,” *26th Canadian Conf. on Computational Geometry (CCCG)*, 2014.
- C-183. M.J. Bannister, W.E. Devanny, D. Eppstein and M.T. Goodrich, “The Galois Complexity of Graph Drawing: Why Numerical Solutions are Ubiquitous for Force-Directed, Spectral, and Circle Packing Drawings,” *22nd Int. Symp. Graph Drawing (GD)*, Springer, LNCS, vol. 8871, 149–161, 2014. (Preliminary version of J-86.)
- C-184. M.J. Alam, D. Eppstein, M.T. Goodrich, S. Kobourov and S. Pupyrev, “Balanced Circle Packings for Planar Graphs,” *22nd Int. Symp. Graph Drawing (GD)*, Springer, LNCS, vol. 8871, 125–136, 2014.
- C-185. M. Bannister, M.T. Goodrich, and P. Sampson, “Force-Directed 3D Arc Diagrams,” *22nd Int. Symp. Graph Drawing (GD)*, Springer, LNCS, vol. 8871, 521–522, 2014.
- C-186. M.T. Goodrich and P. Pszona, “Two-Phase Bicriterion Search for Finding Fast and Efficient Electric Vehicle Routes,” *22nd ACM SIGSPATIAL Int. Conf. on Adv. Geographic Information Systems (GIS)*, 193–202, 2014.
- C-187. M.T. Goodrich and J. Simons, “Data-Oblivious Graph Algorithms in Outsourced External Memory,” *8th Int. Conf. on Combinatorial Optimization and Applications (COCOA)*, LNCS, Vol. 8881, 241–257, 2014.
- C-188. M.T. Goodrich, T. Johnson, M. Torres, “Knuthian Drawings of Series-Parallel Flowcharts,” *23rd Int. Symp. on Graph Drawing and Network Visualization (GD)*, Springer, LNCS, vol. 9411, 556–557, 2015. (See also <http://arxiv.org/abs/1508.03931>.)
- C-189. E. Ghosh, M.T. Goodrich, O. Ohrimenko, and R. Tamassia, “Poster: Zero-Knowledge Authenticated Order Queries and Applications,” *IEEE Symp. on Security and Privacy*, 2015. (See also <https://eprint.iacr.org/2015/283>.)
- C-190. M.T. Goodrich and A. Eldawy, “Parallel Algorithms for Summing Floating-Point Numbers,” *28th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, 13–22, 2016.
- C-191. W.E. Devanny, M.T. Goodrich, and K. Jetviroj, “Parallel Equivalence Class Sorting: Algorithms, Lower Bounds, and Distribution-Based Analysis,” *28th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, 265–274, 2016.
- C-192. D. Eppstein, M.T. Goodrich, J. Lam, N. Mamano, M. Mitzenmacher, and M. Torres, “Models and Algorithms for Graph Watermarking,” *19th Information Security Conf. (ISC)*, 2016. **Best Student Paper Award.**
- C-193. E. Ghosh, M.T. Goodrich, O. Ohrimenko, R. Tamassia, “Verifiable Zero-Knowledge Order Queries and Updates for Fully Dynamic Lists and Trees,” *10th Conf. on Security and Cryptography for Networks (SCN)*, 216–236, 2016.
- C-194. M.T. Goodrich, E. Kornaropoulos, M. Mitzenmacher, R. Tamassia, “More Practical and Secure History-Independent Hash Tables,” *21st European Symp. on Research in Computer Security (ESORICS)*, 2016.
- C-195. J.J. Besa Vial, W.E. Devanny, D. Eppstein, and M.T. Goodrich, “Scheduling Autonomous Vehicle Platoons Through an Unregulated Intersection,” *2016 Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS)*.
- C-196. M.J. Alam, M.B. Dillencourt, and M.T. Goodrich, “Capturing Lombardi Flow in Orthogonal Drawings by Minimizing the Number of Segments,” *24th Int. Symp. on Graph Drawing and Network Visualization (GD)*, 2016.
- C-197. M.J. Alam, M.T. Goodrich, and T. Johnson, “Sibling-First Recursive Graph Drawing for Java Bytecode,” *24th Int. Symp. on Graph Drawing and Network Visualization (GD)*, 2016.



- C-198. M.T. Goodrich, S. Gupta, and M. Torres, “A Topological Algorithm for Determining How Road Networks Evolve Over Time,” *24th ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (GIS)*, 2016.
- C-199. M.J. Alam, M.T. Goodrich, and T. Johnson, “J-Viz: Finding Algorithmic Complexity Attacks via Graph Visualization of Java Bytecode,” *13th IEEE Symp. on Visualization for Cyber Security (VizSec)*, 2016.
- C-200. M.T. Goodrich, E. Kornaropoulos, M. Mitzenmacher, and R. Tamassia, “Auditable Data Structures,” *2nd IEEE Symp. on Security and Privacy (EuroS&P)*, 2017.

#### Other Publications:

- O-1. M.T. Goodrich, “Guest Editor’s Introduction,” *International Journal of Computational Geometry & Applications*, **2**(2), 1992, 113–116.
- O-2. M.T. Goodrich, “Parallel Algorithms Column 1: Models of Computation,” *SIGACT News*, **24**(4), 1993, 16–21.
- O-3. M.T. Goodrich, V. Mirelli, M. Orletsky, and J. Salowe, “Decision tree construction in fixed dimensions: Being global is hard but local greed is good,” Technical Report TR-95-1, Johns Hopkins University, Department of Computer Science, Baltimore, MD 21218, May 1995.
- O-4. R. Tamassia, P.K. Agarwal, N. Amato, D.Z. Chen, D. Dobkin, R.L.S. Drysdale, S. Fortune, M.T. Goodrich, J. Hersberger, J. O’Rourke, F.P. Preparata, J.-R. Sack, S. Suri, I.G. Tollis, J.S. Vitter, and S. Whitesides, “Strategic Directions in Computational Geometry Working Group Report,” *ACM Computing Surveys*, **28A**(4), December 1996.
- O-5. G.A. Gibson, J.S. Vitter, and J. Wilkes, A. Choudhary, P. Corbett, T.H. Cormen, C.S. Ellis, M.T. Goodrich, P. Highnam, D. Kotz, K. Li, R. Muntz, J. Pasquale, M. Satyanarayanan, D.E. Vengroff, “Report of the Working Group on Storage I/O Issues in Large-Scale Computing,” *ACM Computing Surveys*, **28A**(4), December 1996.
- O-6. T.H. Cormen and M.T. Goodrich, “A Bridging Model for Parallel Computation, Communication, and I/O,” *ACM Computing Surveys*, **28A**(4), December 1996.
- O-7. M.T. Goodrich, “Computer Science Issues in the National Virtual Observatory,” in *Virtual Observatories of the Future*, ASP Conf. Series, vol. 225, R.J. Brunner, S.G. Djorgovski, and A.S. Szalay, eds., 329–332, 2001.
- O-8. M.T. Dickerson and M.T. Goodrich, “Matching Points to a Convex Polygonal Boundary,” Proceedings of the 13th Canadian Conf. on Computational Geometry (CCCG’01), 89–92, 2001.
- O-9. M.T. Goodrich, “Guest Editor’s Foreword,” *Algorithmica*, **33**(3), 271, 2002.
- O-10. M.T. Goodrich, M. Shin, C.D. Straub, and R. Tamassia, “Distributed Data Authentication (System Demonstration),” *DARPA Information Survivability Conf. and Exposition*, IEEE Press, Volume 2, 58–59, 2003.
- O-11. M.T. Goodrich and R. Tamassia, “Efficient and Scalable Infrastructure Support for Dynamic Coalitions,” *DARPA Information Survivability Conf. and Exposition*, IEEE Press, Volume 2, 246–251, 2003.

#### PROFESSIONAL SERVICE

*Guest Editor:*

*Int. Journal of Computational Geometry & Applications*, **2**(2), 1992  
*Journal of Computer & System Sciences*, **52**(1), 1996

*Computational Geometry: Theory and Applications*, **12**(1–2), 1999.  
*Algorithmica*, **33**(3), 2002.

*Editorial Board Membership:*

*Computational Geometry: Theory and Applications*, 2006–2015  
*Journal of Computer & System Sciences*, 1994–2011  
*Journal of Graph Algorithms and Applications*, 1996–2011  
*Int. Journal of Computational Geometry & Applications*, 1993–2010  
*Information Processing Letters*, 1995–1997

*Journal Advisory Board Membership:*

*Int. Journal of Computational Geometry & Applications*, 2010–  
*Journal of Graph Algorithms and Applications*, 2011–

*Program Committee Service:*

7th ACM Symp. on Computational Geometry (SoCG), 1991  
1991 Workshop on Algorithms and Data Structures (WADS)  
8th ACM Symp. on Computational Geometry (SoCG), 1992  
25th ACM Symp. on Theory of Computing (STOC), 1993  
*Chair*, 26th ACM Symp. on Theory of Computing (STOC), 1994  
11th ACM Symp. on Computational Geometry (SoCG), 1995  
DAGS '95 Conf. on Electronic Publishing and the Information Superhighway  
1996 SIAM Discrete Mathematics Conference  
1997 Workshop on Algorithms and Data Structures (WADS)  
International Symposium on Graph Drawing (GD), 1997  
1999 Workshop on Algorithms and Data Structures (WADS)  
*Co-chair*, Workshop on Algorithm Engineering and Experimentation (ALENEX), 1999  
International Symposium on Graph Drawing (GD), 2000  
2000 Workshop on Algorithm Engineering (WAE)  
41st IEEE Symp. on Foundations of Computer Science (FOCS), 2000  
2001 Workshop on Algorithms and Data Structures (WADS)  
International Symposium on Graph Drawing (GD), 2001  
Workshop on Algorithm Engineering and Experimentation (ALENEX), 2002  
18th ACM Symp. on Computational Geometry (SoCG), 2002  
13th ACM-SIAM Symp. on Discrete Algorithms (SODA), 2002  
*Co-Chair*, Graph Drawing 2002  
International Symposium on Graph Drawing (GD), 2003  
16th ACM-SIAM Symp. on Discrete Algorithms (SODA), 2005  
32nd Int. Colloq. on Automata, Languages and Programming (ICALP), 2005  
12th Int. Computing and Combinatorics Conference (COCOON), 2006  
13th ACM Conf. on Computer and Communication Security (CCS), 2006  
15th Annual European Symposium on Algorithms (ESA), 2007  
5th International Conference on Applied Cryptography and Network Security (ACNS), 2007  
21st IEEE International Parallel & Distributed Processing Symposium (IPDPS), 2007  
19th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2007  
5th Workshop on Algorithms and Models for the Web-Graph (WAW), 2007  
7th International Workshop on Experimental Algorithms (WEA), 2008  
Second International Frontiers of Algorithmics Workshop (FAW), 2008

16th ACM Int. Symp. on Advances in Geographic Information Systems (GIS), 2008  
 17th ACM Int. Symp. on Advances in Geographic Information Systems (GIS), 2009  
 31st IEEE Symposium on Security and Privacy (SSP), 2010  
 18th Int. Symp. on Graph Drawing (GD), 2010  
 2011 Workshop on Analytic Algorithmics and Combinatorics (ANALCO)  
 8th Workshop on Algorithms and Models for the Web Graph (WAW), 2011  
 19th International Symposium on Graph Drawing (GD), 2011  
 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2012  
 20th European Symposium on Algorithms (ESA), 2012  
 2013 IEEE Int. Conf. on Big Data (BigData), 2013  
 30th IEEE Int. Conf. on Data Engineering (ICDE), 2014  
 21st ACM Conf. on Computer and Communication Security (CCS), 2014  
 Symposium on Algorithms and Data Structures (WADS), 2015  
 ACM Cloud Computing Security Workshop (CCSW), 2015  
 International Symposium on Graph Drawing (GD), 2015  
*co-chair*, 2016 Workshop on Algorithm Engineering and Experiments (ALENEX)  
 2016 Workshop on Massive Data Algorithmics (MASSIVE)  
 2016 Int. Symposium on Algorithms and Computation (ISAAC)  
 29th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2017

*Conference Committee Service:*

Conference chair, 12th ACM Symposium on Computational Geometry, 1996  
 Organizer, 1st CGC Workshop on Computational Geometry, 1996  
 Co-chair, 1999 Dagstuhl Workshop on Computational Geometry, 1999  
 Conference chair, Graph Drawing 2002

*Steering Committee and Executive Committee Service:*

Member at large, ACM SIG on Algorithms & Comp. Theory (SIGACT) Exec. Comm., 1993–97  
 Member, Exec. comm. for 1996 Federated Computing Research Conference (FCRC)  
 co-Founder and member, Steering Comm. for Workshop on Algorithm Engineering  
 and Experimentation (ALENEX), 1999– (chair since 2014)  
 co-Chair, Steering Comm. for ACM Symposium on Computational Geometry, 1999–2001  
 Member, Steering Comm. for Graph Drawing Conference, 2000–2003, 2014–  
 Conference Chair, ACM SIG on Algorithms & Comp. Theory (SIGACT), 2005–2009

*Review Panelist:*

National Science Foundation, 2000–2016

*Ph.D. Advisees:*

Mujtaba Ghouse	“Randomized Parallel Computational Geometry in Theory and Practice,” May 1993.
Paul Tanenbaum	“On Geometric Representations of Partially Ordered Sets,” May 1995 (co-advised with Prof. Edward Scheinerman).
Mark Orletsky	“Practical Methods for Geometric Searching Problems with Experimental Validation,” May 1996.
Kumar Ramaiyer	“Geometric Data Structures and Applications,” Aug. 1996.
Christian Duncan	“Balanced Aspect Ratio Trees,” Aug. 1999.
Christopher Wagner	“Graph Visualization and Network Routing,” Oct. 1999 (co-advised with Prof. Lenore Cowen).
Stephen Kobourov	“Algorithms for Drawing Large Graphs,” May 2000.
Amitabha Bagchi	“Efficient Strategies for Topics in Internet Algorithmics,” Oct. 2002.

Amitabh Chaudhary “Applied Spatial Data Structures for Large Data Sets,” Oct. 2002.  
 Jeremy Yu Meng “Confluent Graph Drawing,” June 2006.  
 Jonathan Zheng Sun “Algorithms for Hierarchical Structures, with Applications to Security and Geometry,” August 2006.  
 Nodari Sitchinava “Parallel External Memory Model—A Parallel Model for Multi-core Architectures,” Sept. 2009.  
 Darren Strash “Algorithms for Sparse Geometric Graphs and Social Networks,” May 2011.  
 Lowell Trott “Geometric Algorithms for Social Network Analysis,” May 2013.  
 Joseph Simons “New Dynamics in Geometric Data Structures,” May 2014.  
 Pawel Pszona “Practical Algorithms for Sparse Graphs,” May 2014.

*Ph.D. Committee Service:*

John Augustine	UC-Irvine	Advancement to candidacy, September 2003
Nikos Triandopoulos	Brown U.	Thesis prelim., February 2004
Einar Mykletun	UC-Irvine	Advancement to candidacy, March 2004
Kartic Subr	UC-Irvine	Advancement to candidacy, September 2004
S. Joshua Swamidass	UC-Irvine	Advancement to candidacy, April 2005
Jeong Hyun Yi	UC-Irvine	Thesis defense, August, 2005
Nodari Sitchinava	UC-Irvine	Advancement to candidacy, chair, December 2005
John Augustine	UC-Irvine	Thesis defense, July 2006
Maithili Narasimha	UC-Irvine	Thesis defense, August, 2006
Josiah Carlson	UC-Irvine	Advancement to candidacy, August 2006
Xiaomin Liu	UC-Irvine	Advancement to candidacy, September 2006
Gabor Madl	UC-Irvine	Advancement to candidacy, September 2006
Nikos Triandopoulos	Brown U.	Thesis defense, September 2006
Rabia Nuray-Turan	UC-Irvine	Advancement to candidacy, May 2007
S. Joshua Swamidass	UC-Irvine	Thesis defense, June 2007
Michael Sirivianos	UC-Irvine	Advancement to candidacy, June 2007
Kevin Wortman	UC-Irvine	Advancement to candidacy, August 2007
Di Ma	UC-Irvine	Advancement to candidacy, December 2007
Josiah Carlson	UC-Irvine	Thesis defense, December 2007
Michael Nelson	UC-Irvine	Advancement to candidacy, chair, March 2008
Minas Gjoka	UC-Irvine	Advancement to candidacy, June 2008
Sara Javanmardi	UC-Irvine	Advancement to candidacy, June 2008
Ali Zandi	UC-Irvine	Advancement to candidacy, September 2008
Jihye Kim	UC-Irvine	Thesis defense, September 2008
Darren Strash	UC-Irvine	Advancement to candidacy, December 2008
Kevin Wortman	UC-Irvine	Topic defense, January 2009
Nodari Sitchinava	UC-Irvine	Topic defense, chair, June 2009
Fabio Soldo	UC-Irvine	Advancement to candidacy, July 2009
Emil De Cristofaro	UC-Irvine	Advancement to candidacy, July 2009
Di Ma	UC-Irvine	Thesis defense, August 2009
Yanbin Lu	UC-Irvine	Advancement to candidacy, December 2009
Anh Le	UC-Irvine	Advancement to candidacy, April 2010
Lowell Trott	UC-Irvine	Advancement to candidacy, June 2010
Xiaomin Liu	UC-Irvine	Thesis defense, August 2010

Josh Olsen	UC-Irvine	Advancement to candidacy, September 2010
Yasser Altowim	UC-Irvine	Advancement to candidacy, December 2010
Angela Wong	UC-Irvine	Advancement to candidacy, May 2011
Joshua Hill	UC-Irvine	Advancement to candidacy, September 2011
Alex Abatzoglou	UC-Irvine	Advancement to candidacy, September 2011
Michael Wolfe	UC-Irvine	Masters Thesis defense, October 2011
Olya Ohrimenko	Brown Univ.	PhD Thesis proposal, October 2011
Yanbin Lu	UC-Irvine	PhD Thesis defense, November 2011
Chun Meng	UC-Irvine	Advancement to candidacy, December 2011
Abinesh Ramakrishnan	UC-Irvine	Advancement to candidacy, March 2012
Pegah Sattari	UC-Irvine	PhD Thesis defense, April 2012
Michael Bannister	UC-Irvine	PhD Thesis defense, May 2015
Yingyi Bu	UC-Irvine	PhD Thesis defense, August 2015
Jenny Lam	UC-Irvine	PhD Thesis defense, November 2015
Timothy Johnson	UC-Irvine	Advancement to candidacy, chair, June 2016
Jiayu Xu	UC-Irvine	Advancement to candidacy, November 2016
Sky Faber	UC-Irvine	PhD Thesis defense, November 2016

*University Service:*

Ph.D. Requirements Committee, Dept. of Computer Science, chair: 1987–89  
Graduate Admissions Committee, Dept. of Computer Science, 1991–1993 (chair: 1992)  
Faculty Recruiting Committee, Dept. of Computer Science, 1993,95,96 (chair: 1996)  
Steering Committee, Whiting School of Engineering, 1990–93 (chair, 1993)  
Johns Hopkins Homewood Academic Computing Oversight Committee, 1990–93  
Curriculum Committee, Whiting School of Engineering, 1994–96  
Strategic Planning Committee, Whiting School of Engineering, 1999–00  
Graduate Policy Committee, UCI Dept. of Information & Computer Science (ICS), 2001–02  
Faculty Search Committee in Cryptography, UCI Dept. of ICS, 2001–03  
School of Info. and Computer Science Executive Committee, 2002–04  
UCI Committee on Educational Policy (CEP), 2002–03, 2004–06  
UCI Change of Major Criteria Committee, 2002–03  
UCI CEP Policy Subcommittee, 2002–2003  
Distinguished Faculty Search Committee, Bren School of ICS, 2004–11 (chair, 2007–08)  
Equity Advisor, Bren School of ICS, 2005–09  
Dean’s Advisory Council, Bren School of ICS, 2007–13  
Associate Dean for Faculty Development, Bren School of ICS, 2006–12  
Chair, Department of Computer Science, Bren School of ICS, 2012–13  
Master of Computer Science Development Committee, Bren School of ICS, 2013–2016  
Stragic Planning Committee, Dept. of Computer Science, Bren School of ICS, 2015–16

*Courses Taught and Developed:*

Advanced Parallel Computing (developed and taught at Hopkins)  
Cyber-Puzzlers (designed and taught at UCI)  
Computer Literacy (taught at Purdue, developed at Hopkins)  
Computer Programming for Scientists and Engineers (taught at Purdue)  
Computer Security Algorithms (developed and taught at UCI)  
Computational Models (revised and taught at Hopkins)

Computational Geometry (revised and taught at Hopkins and UCI)  
 Compiler Theory and Design (revised and taught at Hopkins)  
 Computer Graphics (taught at Hopkins)  
 Cyber-Fraud Detection and Prevention (designed and taught at UCI)  
 Data Structures (revised and taught at Hopkins and UCI)  
 Graph Algorithms (revised and taught at UCI)  
 Formal Languages and Automata Theory (revised and taught at UCI)  
 Introduction to Algorithms (developed and taught at Hopkins and UCI)  
 Internet Algorithmics (developed and taught at Hopkins, Brown, and UCI)  
 Design and Analysis of Algorithms (revised and taught at Hopkins and UCI)  
 Parallel Algorithms (developed and taught at Hopkins and Univ. of Illinois)

*Scientific Consulting:*

APAC Security, Inc., 2005  
 Algomagic Technologies, Inc., 2000–2005  
 Army Research Laboratory, Fort Belvoir, 1995  
 AT&T, 1998  
 Battelle Research Triangle, Columbus Division, 1996  
 Brown University, 2000–2007  
 3M, 2015  
 Purdue University, 2002  
 The National Science Foundation, 1990–2016  
 Univ. of Miami, 1999  
 Walt Disney Animation Studios, 2009

*Technical Expert Consulting:*

- 2011–12, Technical expert, Sidley Austin LLP (Los Angeles) and IGT Technologies
- 2012, Technical expert and deponent, Jones Day (Irvine) and Quiksilver
- 2013–14, Technical expert and deponent, Kirkland & Ellis, LLP (New York) and Content-Guard Holdings
- 2013–14, Technical expert, Kirkland & Ellis, LLP (Chicago) and Apple
- 2013–14, Technical expert and deponent, Jeffer Mangels Butler & Mitchell, LLP (Irvine) and Uniloc
- 2014–15, Technical expert, deponent, and testifying witness, McKool Smith, PC (Dallas) and ContentGuard Holdings
- 2014–16, Technical expert and deponent, Fenwick & West, LLP (San Francisco) and Symantec
- 2014–, Technical expert, Kirkland & Ellis, LLP (Chicago)
- 2016, Technical expert, deponent, and testifying witness, Dentons US LLP
- 2016–, Technical expert, Kramer Levin Naftalis & Frankel LLP (Menlo Park, CA) and Acceleration Bay LLC
- 2016–, Technical expert, Sheppard Mullin Richter & Hampton LLP (Los Angeles)
- 2016–, Technical expert, Kramer Levin Naftalis & Frankel LLP (Menlo Park, CA) and Finjan, Inc.

**GRANTS AND CONTRACTS**

1. PI, "Research Initiation Award: Parallel and Sequential Computational Geometry," National Science Foundation (NSF Grant CCR-8810568), \$32,914, 1988–90.
2. co-PI, "Paradigms for Parallel Algorithm Design," NSF and DARPA (as NSF Grant CCR-8908092), \$523,837, 1989–93 (with S.R. Kosaraju (PI), S. Kasif, and G. Sullivan).
3. PI, "Parallel Computation and Computational Geometry," NSF (Grant CCR-9003299), \$67,436, 1990–93.
4. co-PI, "A Facility for Experimental Validation," NSF (Grant CDA-9015667), \$1,476,147, 1991–96 (with G. Masson (PI), J. Johnstone, S. Kasif, S.R. Kosaraju, S. Salzberg, S. Smith, G. Sullivan, L. Wolff, and A. Zwarico).
5. PI, "Parallel Network Algorithms for Cell Suppression," The Bureau of the Census (JSA 91-23), \$14,998 1991–92.
6. PI, "A Geometric Framework for the Exploration & Analysis of Astrophysical Data," NSF (Grant IRI-9116843), \$535,553, 1991–96 (with S. Salzberg and H. Ford (from Physics and Astronomy Dept.)).
7. PI, "Research Experiences for Undergraduates supplement to IRI-9116843," NSF, \$4,000, 1993–94 (with S. Salzberg and H. Ford).
8. PI, "Constructing, Maintaining, and Searching Geometric Structures," NSF (Grant CCR-9300079), \$134,976, 1993–96.
9. co-PI, "Robust and Applicable Geometric Computing," Army Research Office (ARO MURI Grant DAAH04-96-1-0013), \$4,500,000, 1996–2000 (with F. Preparata (PI, Brown U.), R. Tamassia (Brown U.), S. Rao Kosaraju, J. Vitter (Duke U.), and P. Agarwal (Duke U.)). Subaward size: \$1,466,640.
10. PI, "Application-Motivated Geometric Algorithm Design," NSF (Grant CCR-9625289), \$107,389, 1996–98.
11. co-PI, "vBNS Connectivity for the Johns Hopkins University," NSF, \$350,000, 1997–99 (with T.O. Poehler (PI), D.J. Binko, J.G. Neal, and A.S. Szalay).
12. co-PI, "Product Donation, Technology for Education Program," Intel Corporation, \$480,071, 1997–2001 (with T.O. Poehler (PI), J.H. Anderson, A.S. Szalay, and M. Robbins).
13. co-PI, "A Networked Computing Environment for the Manipulation & Visualization of Geometric Data" (Research Infrastructure), NSF, \$1,638,785, 1997–2003 (with L.B. Wolff (PI), Y. Amir, S.R. Kosaraju, S. Kumar, R. Tamassia (Brown U.), R.H. Taylor, and D. Yarowsky).
14. PI, "Geometric Algorithm Design and Implementation," NSF, Grant CCR-9732300, \$224,982, 1998–2002.
15. PI, "Certification Management Infrastructure – Certificate Revocation," \$52,023, 1998, NSA LUCITE grant.
16. PI, "Software Engineering Data Loading, Analysis, and Reporting," \$41,614, 1998, NSA LUCITE grant.
17. PI, "Establishing a LUCITE Collaboration Environment," \$10,018, 1998, NSA LUCITE grant.
18. PI, "In Support of a Secure Multilingual Collaborative Computing Environment," \$51,471, 1999–2000, NSA LUCITE grant.
19. PI, "Accessing Large Distributed Archives in Astronomy and Particle Physics," \$199,981. subcontract to UCI from Johns Hopkins Univ. on NSF Grant PHY-9980044 (total budget, \$2,500,000), 1999–2004.

20. PI, "Efficient and Scalable Infrastructure Support for Dynamic Coalitions," \$1,495,000, DARPA Grant F30602-00-2-0509, 2000-2003 (with Robert Cohen and Roberto Tamassia), including \$227,893 subaward to UCI (with Gene Tsudik).
21. PI, "Graph Visualization and Geometric Algorithm Design," \$400,000, NSF Grant CCR-0098068, 2001-2004 (with Roberto Tamassia).
22. PI, "Collaborative Research: Teaching Data Structures to the Millennium Generation," \$125,000, NSF Grant DUE-0231467, 2003-2005.
23. PI, "Collaborative Research: An Algorithmic Approach to Cyber-Security," \$100,000, NSF Grant CCR-0311720, 2003-2006.
24. PI, "The OptIPuter," \$900,000, subcontract from UCSD on NSF ITR grant CCR-0225642 (total budget, \$13.5 million), 2002-2007 (with Padhraic Smyth and Kane Kim).
25. PI, "ITR: Algorithms for the Technology of Trust," \$300,000, NSF Grant CCR-0312760, 2003-2009.
26. co-PI, "SDCI Data New: Trust Management for Open Collaborative Information Repositories: The CalSWIM Cyberinfrastructure," NSF grant OCI-0724806, \$1,103,590, 2007-2012.
27. co-PI, "Support for Machine Learning Techniques for Cyber-Fraud Detection," Experian Corporation, \$200,000 gift, 2008.
28. PI, "IPS: Collaborative Research: Privacy Management, Measurement, and Visualization in Distributed Environments," NSF Grant IIS-0713046, \$224,851, 2007-2009.
29. PI, "Collaborative Research: Algorithms for Graphs on Surfaces," \$400,000, NSF Grant CCR-0830403, 2008-2011.
30. PI, "ROA Supplement: IPS: Collaborative Research: Privacy Management, Measurement, and Visualization in Distributed Environments," NSF Grant IIS-0847968, \$25,000, 2008-2009.
31. co-investigator, "Scalable Methods for the Analysis of Network-Based Data," Office of Naval Research: Multidisciplinary University Research Initiative (MURI) Award, number N00014-08-1-1015, \$529,152, 2008-2014.
32. PI, "EAGER: Usable Location Privacy for Mobile Devices," NSF Grant 0953071, \$300,000, 2009-2011.
33. PI, "TC:Large:Collaborative Research: Towards Trustworthy Interactions in the Cloud," NSF Grant 1011840, \$500,000, 2010-2015.
34. PI, "TWC: Medium: Collaborative: Privacy-Preserving Distributed Storage and Computation," NSF Grant 1228639, \$390,738, 2012-2016.
35. PI, "Support for Research on Geometric Motion Planning," 3M Corporation, \$40,000 gift, 2014.
36. PI, "A4V: Automated Analysis of Algorithm Attack Vulnerabilities," subcontract 10036982-UCI from University of Utah for DARPA agreement no. AFRL FA8750-15-2-0092, \$980,000, 2015-2019.
37. PI, "TWC: Small: Collaborative: Practical Security Protocols via Advanced Data Structures," NSF Grant 1526631, \$166,638, 2015-2017.

#### **SELECTED INVITED TALKS (RECENT YEARS ONLY)**

- "Probabilistic Packet Marking for Large-Scale IP Traceback," Purdue Univ., 2003
- "Algorithms for Data Authentication," Harvey Mudd College, 2003
- "Efficient Tree-Based Revocation in Groups of Low-State Devices," Univ. of Arizona, 2004



- “Leap-Frog Packet Linking and Diverse Key Distributions for Improved Integrity in Network Broadcasts,” Southern California Security and Cryptography Workshop, 2005
- “Is Your Business Privacy Protected?,” NEXT Connections, 2005
- “Distributed Peer-to-peer Data Structures,” Harvard Univ., 2006
- “Balancing Life with an Academic Research Career,” Grace Hopper Conference, 2006
- “Computer Security in the Large,” Univ. Texas, San Antonio, 2006
- “Inspirations in Parallelism and Computational Geometry,” Brown Univ., 2006
- “Efficiency and Security Issues for Distributed Data Structures,” Computer Science Distinguished Lecture Series, Johns Hopkins Univ., 2006
- “Efficiency and Security Issues for Distributed Data Structures,” UCLA, 2006
- “Efficiency and Security Issues for Distributed Data Structures,” Edison Distinguished Lecturer Series, Univ. of Notre Dame, 2006
- “Efficiency and Security Issues for Distributed Data Structures,” Computer Science Distinguished Lecturer Series, Texas A & M Univ., 2006
- “Algorithms for Secure Computing and Searching with Applications to Medical Informatics,” Purdue Univ., 2006
- “Blood on the Computer: How Algorithms for Testing Blood Samples can be Used for DNA Sequencing, Wireless Broadcasting, and Network Security,” Univ. of Southern California, 2007
- “Blood on the Computer: How Algorithms for Testing Blood Samples can be Used for DNA Sequencing, Wireless Broadcasting, and Network Security,” Univ. California, San Diego, 2007
- “Blood on the Computer: How Algorithms for Testing Blood Samples can be Used for DNA Sequencing, Wireless Broadcasting, and Network Security,” Univ. Minnesota, 2007
- “Blood on the Database: How Algorithms for Testing Blood Samples can be Used for Database Integrity,” Invited Keynote, 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security (DBSec), 2007
- “Space-Efficient Straggler Identification,” ALCOM Seminar, Univ. of Aarhus, 2007
- “Blood on the Computer: How Algorithms for Testing Blood Samples can be used in Modern Applications,” ALCOM Seminar, Univ. of Aarhus, 2007
- “Studying Road Networks Through an Algorithmic Lens,” ALCOM Seminar, Univ. of Aarhus, 2008
- “Studying Geometric Graph Properties of Road Networks Through an Algorithmic Lens,” International Workshop on Computing: from Theory to Practice, 2009
- “Randomized Shellsort: A Simple Oblivious Sorting Algorithm,” Distinguished Lecture Series, Department of Computer Science, Brown University, 2009
- “Simulating Parallel Algorithms in the MapReduce Framework with Applications to Parallel Computational Geometry,” MASSIVE 2010
- “Data Cloning Attacks for Nearest-Neighbor Searching based on Retroactive Data Structures,” Department of Computer Science, UCSB, 2011
- “Turning Privacy Leaks into Floods: Surreptitious Discovery of Social Network Friendships and Other Sensitive Binary Attribute Vectors,” Department of Computer Science Distinguished Lecturer Series, Univ. of Illinois, Chicago, 2011
- “Turning Privacy Leaks into Floods: Surreptitious Discovery of Social Network Friendships and Other Sensitive Binary Attribute Vectors,” Department of Computer Sciences, Purdue

Univ., 2011

- “Spin-the-bottle Sort and Annealing Sort: Oblivious Sorting via Round-robin Random Comparisons,” Department of Computer Science, Brown Univ., 2012
- “Using Data-Oblivious Algorithms for Private Cloud Storage Access,” Qatar University, 2013
- “Using Data-Oblivious Algorithms for Private Cloud Storage Access,” Department of Computer Science and Engineering Distinguished Speaker Series, University of Buffalo, 2013
- “Force-Directed Graph Drawing Using Social Gravity and Scaling,” invited talk, ICERM Workshop on Stochastic Graph Models, Providence, RI, 2014
- “Invertible Bloom Lookup Tables and Their Applications in Large-Scale Data Analysis,” invited key-note speaker, Algorithms for Big Data, Frankfurt, Germany, 2014
- “Invertible Bloom Lookup Tables and Their Applications in Large-Scale Data Analysis,” Brown University, Providence, RI, 2014
- “Studying Road Networks Through an Algorithmic Lens,” Bold Aspirations Visitor and Lecture, University of Kansas, 2015
- “Learning Character Strings via Mastermind Queries, with Case Studies,” Invited Lecture, Workshop on Pattern Matching, Data Structures and Compression, Bar-Ilan University, Tel Aviv, Israel, 2016
- “Invertible Bloom Lookup Tables and Their Applications in Data Analysis,” University of Hawaii, 2016
- “Invertible Bloom Lookup Tables,” Purdue University, 2016
- “Combinatorial Pair Testing: Distinguishing Workers from Slackers,” Calvin College, 2016
- “Invertible Bloom Lookup Tables,” University of California, Riverside, 2016

#### **PROFESSIONAL SOCIETY MEMBERSHIPS**

American Association for the Advancement of Science (AAAS), Fellow

IEEE and IEEE Computer Society, Fellow

Association for Computing Machinery (ACM), Fellow