IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent of: Baek et al.	§	Petition for *Inter Partes* Review
	§	
U.S. Patent No. 6,978,346	§	Attorney Docket No.: 47415.430
	§	
Issued: December 20, 2005	§	Customer No.: ______
	§	
Title: APPARATUS FOR	§	Real Parties in Interest: Dell Inc.,
REDUNDANT INTER-	§	Hewlett-Packard Company, and NetApp,
CONNECTION	§	Inc.
BETWEEN MULTIPLE	§	
HOSTS AND RAID	§	

Declaration of Dr. M. Ray Mercer

Under 37 C.F.R. § 1.68

I, Dr. M. Ray Mercer, do hereby declare:

1. I am making this declaration at the request of Dell Inc., Hewlett-Packard Company, and NetApp, Inc. in the matter of the *Inter Partes* Review of U.S. Patent No 6,978,346 ("the '346 Patent") to Baek et al.

2. I am being compensated for my work in this matter. My compensation in no way depends upon the outcome of this proceeding.

3. In the preparation of this declaration, I have studied:

(1) The '346 patent, DHPN-1001;

(2) The prosecution history of the '346 patent, DHPN-1002;

(3) Peter Weygant, *Clusters for High Availability: A Primer of HP-UX Solutions*, 1996 ("Weygant"), DHPN-1003;

(4) *Managing MC/ServiceGuard*, Hewlett-Packard Company, 1998 ("ServiceGuard"), DHPN-1004;

(5) Hathorn et al., U.S. Pat. No. 5,574,950 ("the '950 patent"), DHPN-1005;

(6) Surugguchi et al., International Publication No. WO 99/38067 ("Mylex"), DHPN-1007; and

(7) *American National Standard for Information Technology – Fibre Channel Arbitrated Loop (FC-AL-2)*, June 28, 1999 ("ANSI"), DHCP-1008.

4. In forming the opinions expressed below, I have considered:

(1) The documents listed above,

(2) The relevant legal standards, including the standard for obviousness provided in *KSR International Co. v. Teleflex, Inc.*, 550 U.S. 398 (2007) and any additional authoritative documents as cited in the body of this declaration, and

(3) My knowledge and experience based upon my work in this area as described below.

Qualifications and Professional Experience

5. My qualifications are set forth in my curriculum vitae, a copy of which is attached as Appendix 1.

6. I have over 45 years of dual industrial and academic experience in Electrical Engineering and Computer Engineering. I received a B.S. in Electrical Engineering from Texas Tech University in 1968. From 1968 to 1973, I was a Research/Development Engineer at General Telephone and Electronics Sylvania in Mountain View, California, and I received an M.S. in Electrical Engineering from Stanford University in 1971. From 1973 to 1977, I was a Member of Technical Staff at Hewlett-Packard's Santa Clara Division and subsequently at Hewlett-Packard Laboratories in Palo Alto, California. From 1977 to 1980, I was a Lecturer in the Division of Mathematics, Statistics, and Computer Science at the University of Texas at San Antonio, and I received a Ph.D. in Electrical Engineering from the University of Texas at Austin in 1980. From 1980 to 1983, I was a Member of Technical Staff at Bell Laboratories in Murray Hill, New Jersey.

7. In 1983, I was appointed Assistant Professor of Electrical and Computer Engineering at the University of Texas at Austin. In 1987, I was promoted to Associate Professor and in 1991, Professor. In 1995, I was appointed

Professor of Electrical and Computer Engineering, Leader of the Computer Engineering Group and Holder of the Computer Engineering Chair at Texas A&M University in College Station, Texas. My teaching, my research, my technical publications, and my supervision of graduate students during this period included the areas of computer clusters, redundant connections, and networking – key issues in this proceeding.

8. In September 2005, I retired, and the Regents of the Texas A&M University System appointed me as Professor Emeritus of Electrical and Computer Engineering at Texas A&M University.

9. Since 1984, I have been an independent consultant and provided private consultation and advice in Electrical and Computer Engineering to numerous entities including IBM, Inc., Rockwell International, Motorola Semiconductor, AT&T, Inc., and SigmaTel. I also have been hired by numerous law firms to provide them and their clients with expert consultation and expert testimony – often in the areas of patent infringement litigation related to Electrical and Computer Engineering.

10. I was actively involved in numerous professional organizations including the Institute of Electrical and Electronics Engineers ("IEEE"), and I was recognized as an IEEE Fellow in 1994. I was the Program Chairman for the 1989 International Test Conference, which is an IEEE-sponsored annual conference with (at that time) more than one thousand attendees and over one hundred presented papers. I won the Best Paper Award at the 1982 International Test Conference. I also won a Best Paper Award at the 1991 Design Automation Conference, an annual conference with (at that time) more than ten thousand attendees and five hundred submitted papers, many of which related to the design of integrated circuit based systems. The subject of this paper involved trade-offs between power consumption and processing speed in integrated circuits. I also won a Best Paper Award at the 1999 VLSI Test Symposium. I am the inventor on United States patents that relate to the design of integrated circuits. I was selected as a National Science Foundation Presidential Young Investigator in 1986.

11. I am familiar with the knowledge and capabilities one of ordinary skill in the networking and computing cluster arts in the period around 2000. Specifically, my work with students, undergraduates as well as masters and Ph.D. candidates, with colleagues in academia, and with engineers practicing in industry allowed me to become personally familiar with the level of skill of individuals and the general state of the art. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the networking and computing cluster arts in the period around 2000 – the period that includes the filing date of the '346 patent.

Relevant Legal Standards

12. I have been asked to provide my opinions regarding whether the claims of the '346 patent are anticipated or would have been obvious to a person having ordinary skill in the art at the time of the alleged invention, in light of the prior art. It is my understanding that, to anticipate a claim under 35 U.S.C. § 102, a reference must teach every element of the claim. Further, it is my understanding that a claimed invention is unpatentable under 35 U.S.C. § 103 if the differences between the invention 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 to which the subject matter pertains. I also understand that the obviousness analysis takes into account factual inquiries including the level of ordinary skill in the art, the scope and content of the prior art, and the differences between the prior art and the claimed subject matter.

13. It is my understanding that the Supreme Court has recognized several rationales for combining references or modifying a reference to show obviousness of claimed subject matter. Some of these rationales include the following: combining prior art elements according to known methods to yield predictable results; simple substitution of one known element for another to obtain predictable results; use of a known technique to improve a similar device (method, or product) in the same way; applying a known technique to a known device (method, or product) ready for improvement to yield predictable results; choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; and some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

Background Of '346 patent

14. The '346 Patent relates to a system having "redundant interconnections between multiple hosts and a RAID." Fig. 4 of the '346 patent is especially illustrative and is reproduced below for reference:

FIG. 4

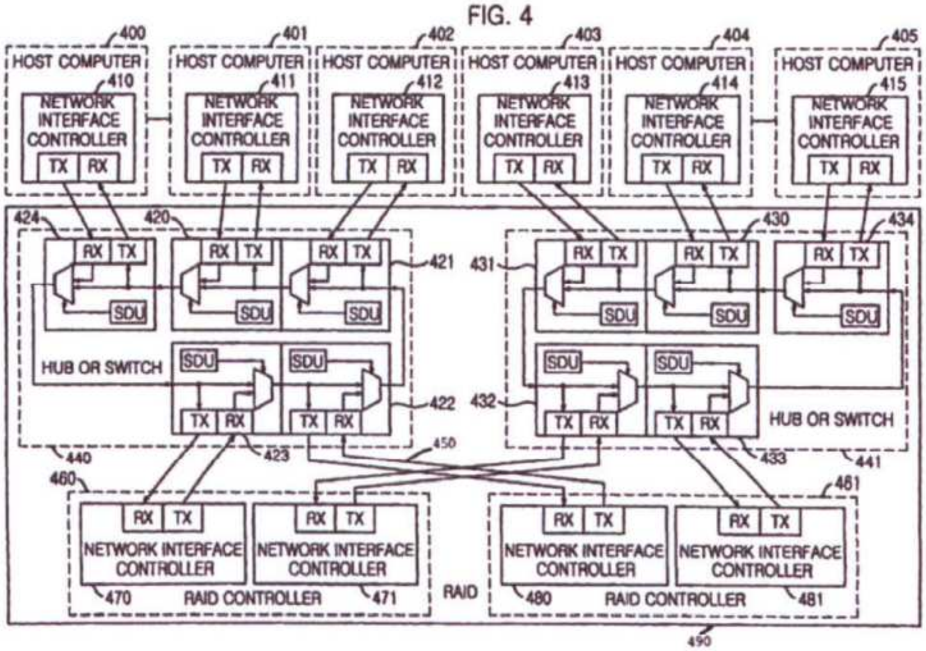

'346 patent, Fig. 4

15. The storage system includes two RAID controllers—460 and 461. Each RAID controller 460, 461 has two Network Interface Controllers (NICs), so

RAID controller 460 includes NICs 470 and 471, and RAID controller 461 includes NICs 480, 481. The system also has two "hub or switch" devices—440 and 441. Each RAID controller is connected to each "hub or switch" device by one of its NICs. RAID controller 460, on the left, is connected to "hub or switch" 440 by NIC 470 and to "hub or switch" 441 by NIC 471. Similarly, RAID controller 461, on the right, is connected to "hub or switch" 441 by NIC 481 and to "hub or switch" 440 by NIC 480.

16. The structure described above provides for a "communication passage between two RAID controllers." '346 Patent, 3:64-65. For instance, a communication passage exists between the RAID controller 460, on the left, and the RAID controller 461, on the right, via NIC 471, switch/hub 441, and NIC 481 (at RAID controller 461). '346 Patent, 3:66 – 4:2. In the same way, a communication passage exists between NIC 481 and NIC 471. '346 Patent, 3:64 – 4:12. Also, a communication passage exists between RAID controller 460, on the left, and RAID controller 461, on the right, via NIC 470, "hub or switch" 440, and NIC 480. *Id.* In the same way, a communication passage exists between NIC 480 and NIC 470. *Id.*

17. The '346 patent fails to provide any examples regarding the types of information that maybe exchanged between the NICs nor any examples regarding how communication paths between the NICs might be used.

18. The system attempts to provide a "fault tolerant function." '346 Patent, 3:63-66. A RAID controller "having [an] error occurrence is removed from the network," then a NIC from other RAID controller "takes over a function" of a NIC on the RAID controller with the error. '346 Patent, 4:19-25. However, such limitation is not reflected in every claim of the '346 patent.

19. Claim 1 provides a basic overview of the teachings of the '346 patent:

> 1. An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:
>
> a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including a first network controlling unit and a second network controlling unit, and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit; and
>
> a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers, wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units, and the first network controlling unit exchanges information with the fourth network controlling unit, and the second network controlling unit exchanges information with the third network controlling unit.

<u>Claim Construction</u>

20. It is my understanding that in order to properly evaluate the '346 patent, the terms of the claims must first be interpreted. It is my understanding that the claims are to be given their broadest reasonable interpretation in light of the specification. It is my further understanding that claim terms are given their ordinary and accustomed meaning as would be understood by one of ordinary skill in the art, unless the inventor, as a lexicographer, has set forth a special meaning for a term.

21. In order to construe the claims, I have reviewed the entirety of the '346 patent, as well as its prosecution history.

network controlling unit, network interface controlling unit

22. These terms appear in claim 1 and in various dependent claims. The specification of the '346 patent does not use the term ***network controlling unit*** or ***network interface controlling*** unit. The terms appear to rely on disclosure in the specification regarding "network interface controllers" for enablement and description. The following passage is an example.

> Network interface controllers, 410 to 415, contained into the host computers, 400 to 405, and the network interface controllers 470, 471, 480, 481 of the RAID controllers 460, 461 are connected with one another by two networks through two hubs 440, 441, and according to a sort of the networks, the network interface controller becomes a fibre channel controller, an ATM controller

and an InfiniBand controller etc. '346 Patent, 3:31-37.

Furthermore, it appears that the claims use the two terms interchangeably. For instance, claim 1 uses ***network controlling unit***, while claim 4 (depending from claim 1) uses the term ***network interface controlling unit***.

23. It is my opinion that a person of ordinary skill in the art would understand the broadest reasonable interpretation of ***network controlling unit*** and ***network interface controlling unit*** in view of the specification to be any component allowing a device to communicate over a network (e.g., Fibre Channel, ATM, or other networks). Furthermore, because of the way the two terms are used interchangeably within the claims, a person of ordinary skill in the art would understand that both terms are intended to mean the same thing.

Network interface controller

24. The specification of the '346 patent uses the term ***network interface controller*** throughout, but does not seek to define or limit the term. See, e.g., '346 Patent, 3:31-37. Furthermore, it appears that claim 9 uses ***network interface controller*** interchangeably with ***network controlling unit***. See, e.g., '346 Patent, 6:31 and 53 (using ***first network controlling unit*** to refer back to ***first network interface controller***).

25. It is my opinion that a person of ordinary skill in the art would understand the broadest reasonable interpretation of ***network interface controller*** in view of the specification to refer to any component allowing a device to communicate over a network (e.g., Fibre Channel, ATM, or other network). Specifically, with no further direction from the specification or the claims, a person of ordinary skill in the art would read the term ***network interface controller*** to be the same as ***network controlling unit*** and ***network interface controlling unit*** (immediately above).

the second network interface controlling unit and the fourth network controlling unit are used for executing a function of the first network interface controlling unit and the third network controlling unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty

26. A literal interpretation of this element from claim 4 is not supported by the specification. For instance, there is no described embodiment in which both the second and fourth network controlling units execute a function of both the first and third network controlling units when a single RAID controller fails. Furthermore, a literal reading of this element does not make sense when the context of claim 1 is taken into account. Specifically, the first and second network controlling units are both on one RAID controller, and the third and fourth network controlling units are both on another RAID controller, according to claim 1. Thus, if the first RAID controller is faulty, the second network controlling unit would not be used, and if the second RAID controller is faulty, the fourth network controlling unit would not be used. With these concerns in mind, a person of ordinary skill in the art would avoid a literal reading of this element.

27. Instead, the specification of the '346 patent states:

> If any one out of two RAID controllers 460, 461 has an occurrence of an error, the RAID controller having the error occurrence is removed from the network, and a second network interface controller of an opposite RAID controller not having the error occurrence takes over a function of a first network interface controller of the RAID controller having the error occurrence. '346 Patent, 4:19-24.

28. It is my opinion that a person of ordinary skill in the art, when grappling with these difficult issues with the literal wording, would understand the broadest reasonable interpretation of the above-recited term to be "if either one of the first RAID controlling unit or second RAID controlling unit has an occurrence of an error, the apparatus uses a network controlling unit of the RAID controlling unit not having the error occurrence." A person of ordinary skill in the art would favor this interpretation because it is consistent with the specification at column 4, lines 19-24 of the '346 patent.

[X] of the at least [Y] connection ports is [are] coupled to one of the first network interface controlling unit and the third network controlling unit

29. In the above-recited feature of claims 5, 6, and 7, ***X*** (***two*** or ***four)*** is the subject, so that X connection ports are coupled as claimed. The term ***one*** is the object of the preposition of the term ***coupled to***, so that ***one of*** the set (where the set is defined as ***the first network interface controlling unit and the third network controlling unit***) is referred to by ***coupled to***. Therefore, a literal and grammatical reading of the above-quoted portion of claims 5, 6, and 7 means that X connecting ports must be coupled to the first network controlling unit or X connecting ports must be coupled to the third network controlling unit, where either condition would satisfy the claim limitation. (Also, see my construction of ***coupled to*** herein below.) However, upon reading the specification at 3:43-47, I believe that the patentee probably intended to say "a connection port is coupled to the first network controlling unit, and another connection port is coupled to the third network controlling unit," in the case of claims 5 and 6 (e.g., port 423 coupled to NIC 470 and port 422 coupled to NIC 480 of Fig. 4). This is a non-literal reading of the claim feature because it is not consistent with a grammatically correct reading of the limitation. In the case of claim 7, which recites ***four*** instead of ***two***, the patentee probably intended to say "out of a total of four connection ports, some of those four connection ports are coupled to the first network controlling unit, and the others of the four connection ports are coupled to the third network controlling unit." However, I do not think that the patent supports such features, as the '346 specification refers to the items 420-424 and 430-434 as "ports," yet there is no disclosed embodiment where more than one of the ports is shown with multiple lines to any one of the NICs in Fig. 4.

30. In the case of claims 5-7, this is evidence that the term ***coupled to*** is broader than "connected to" and, in the context of a hub or switch, ***coupled to*** means that any connection port in a hub or a switch is connected to any other port in a hub or a switch by virtue of the internal structure of the hub or switch. Such a reading of ***coupled to*** would mean that any one of ports 420-424 is coupled to NIC 470 and NIC 480, and any port 430-434 is coupled to NIC 481 and NIC 471 in Fig. 4 directly or indirectly by virtue of the structure of the switch or hub. Because of the above-described tension between the literal, grammatical reading of the phrases and the disclosure in the specification, it is my opinion that a person of ordinary skill in the art, when grappling with these errors in the literal wording of each limitation, if motivated to preserve the validity of a claim, would interpret the broadest reasonable interpretation of ***[X] of the at least [Y] connection ports is [are] coupled to one of the first network interface controlling unit and the third network controlling unit*** to include both a scenario where [X] connection ports are coupled to the same network controlling unit and the scenario where some of the [X] connection ports are coupled to one network interface controlling unit and others of the [X] connection ports are coupled to the other network interface controlling unit: "a hub (or switch) that has at least [Y] ports where at least [X] of the ports are connected directly or indirectly with the first network interface controlling unit or the third network controlling unit."

the rest of the connection ports being provided as a [hub equipment, network switch equipment, switch] connected with the numerous host computers

31. In the above-recited feature of claims 5, 6, and 7, ***connected with*** modifies ***hub equipment***, ***network switch equipment***, or ***switch*** and does not modify ***the connection ports*** simply as a matter of grammar because ***connected with*** immediately follows ***hub equipment***, ***network switch equipment***, and ***switch***. The passage in the '346 specification at 3:48-50 uses the term "the rest," but it does not address ***hub equipment***, ***network switch equipment***, or ***switch*** and is, thus, less illuminating than the grammatical structure of the claim itself. I note that the construction I propose below is not inconsistent with the specification in any event. It is my opinion that a person of ordinary skill in the art would recognize that such construction is consistent with Figs. 4, 5, and 6 of the '346 patent showing a hub or switch connected with the host computers. Furthermore, a person of ordinary skill in the art would recognize that the term ***the rest*** does not exclude that the other ports, coupled to the network controlling units, are also provided as part of the ***hub equipment***, ***network switch equipment***, or ***switch***.

hub

32. The term ***hub*** encompasses both hubs and switches because the '346 patent defines the term as such.

> Herewith, the hubs 440, 441 are provided to connect a system connected to these hubs by one network … and it can be as a hub or a switch. **Hereinafter, they are named a "hub" altogether**. '346 Patent, 3: 13-18 (emphasis added).

Thus, in order to comport with the definition in the specification, the term ***hub*** should be construed as "hub or switch" in its broadest reasonable interpretation.

coupled to

33. This phrase appears only in claims 3, 5, 6, and 7. It does not appear in the specification.

34. The phrase "connected to" appears in claims 2 and 8. As one example, claim 8 contains the phrase "wherein the first network interface controlling unit of the first RAID controlling unit being *connected to* a first connecting unit." (emphasis added)

35. In addition, "connected to" appears in the specification in numerous places. Three examples of the use of this phrase in the specification are cited below:

> Meanwhile, two network interface controllers 470, 471 of the first RAID controller 460 are respectively *connected to* two different hub ports 423,

> 432, and two network interface controllers 480, 481 of the second RAID controller 461 are respectively *connected to* two different hub ports 422, 433. The rest ports 420, 421, 424, 430, 431, 434 of the hubs 440, 441 are *connected to* the host computers 400 to 405. '346 Patent, col. 3, ll. 43-49 (emphasis added).

36. It is my opinion that a person of ordinary skill in the art would understand the broadest reasonable interpretation of ***"coupled to"*** to be broader than the phrase "connected to." For example if entity A is ***"coupled to"*** entity B, then entity A is connected, directly or indirectly, in order to enable the transfer of signals between entities A and B. See also my explanation of the term ***[X] of the at least [Y] connection ports is [are] coupled to one of the first network interface controlling unit and the third network controlling unit***, given above.

<u>*host computers*</u>

37. This term appears in both independent claims 1 and 9 and appears many times in the specification of the '346 patent, e.g., at 3:32 (describing ***host computers*** 400-405 of Figure 4). However, the term is not used in a manner that defines the term nor narrows the term, nor does the specification even appear to give an example of operation of the ***host computers***. Claims 1 and 9 use the term "host computers" (e.g., claim 1-" wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers", claim 9-" wherein the first network interface controller in the first

RAID controller supplies data to the host computers") in the context of the host computers being in communication with the RAID controllers [RAID controlling units].

38. An example of a definition from a technical dictionary from the time is that found in IEEE 100 The Authoritative Dictionary of IEEE standard terms, 7th ed., 2000 ("Host Computer (1): A computer, attached to a network, providing primarily services such as computation, database access or specific programs of special programming languages."), indicating that a host computer is a network computer. It is my opinion that a person of ordinary skill in the art at the time would have understood the broadest reasonable interpretation of ***host computers***, in light of the present specification, to refer to "network connected computers."

RAID controlling unit

39. This phrase appears in claim 1 and its dependent claims, as well as in claim 9. It does not appear in the specification. The term appears to rely on disclosure in the specification regarding "RAID controller" for enablement and description. However, the specification of the '346 patent does not define nor narrow "RAID controller." The following passage is an example:

> As shown in FIG. 4, in the inventive host interface system, a communication circuit is provided in order for an error recovery between two RAID controllers 460, 461, and the bandwidth between

two groups as the host computers 400 to 405 and two RAID controllers 460, 461 becomes twice the single connection bandwidth. Also, in the inventive host interface system, even though one RAID controller 460 or 461 has an occurrence of a trouble, the bandwidth becomes twice the single connection bandwidth. '346 Patent, Col. 3, ll.1-9)

40. It is my opinion that a person of ordinary skill in the art at the time would have understood the broadest reasonable interpretation of ***RAID controlling unit***, in light of the present specification, to refer to "a functional component including hardware that may be controlled by computer code, the functional component providing control to implement RAID storage in an array of storage drives."

RAID controller

41. This term appears in claim 9. As I mentioned above, it is used in the specification, though neither defined nor narrowed. It should also be noted that ***RAID controller*** is used interchangeably with ***RAID controlling unit*** in claim 9 (see, e.g., 6: 35-36—"second RAID controller"—and 6: 55-56—"second RAID controlling unit"). Furthermore, the file history shows at least one place where ***RAID controller*** and ***RAID controlling unit*** were used interchangeably by the applicant. See, e.g., Response to Office Action, Filed August 19, 2004, at the paragraph spanning pages 8-9 (paragraph uses both terms and makes no distinction therebetween). With that in mind, it is my opinion that a person of ordinary skill in the art at the time would have understood the broadest reasonable interpretation of ***RAID controller***, in light of the present specification, to be the same as ***RAID controlling unit*** (immediately above, "a functional component including hardware that may be controlled by computer code, the functional component providing control to implement RAID storage in an array of storage drives").

RAID

42. This term appears in the preambles of claims 1 and 9. The term is also used in claim 9—"a plurality of connection units for connecting the host computers and the RAID; a first and a second RAID controllers, included in the RAID." '346 Patent at 6:23-26. *RAID* is used in the specification to, e.g., refer to RAID 490 of Fig. 4. There is no one definition of the term that is agreed upon by everyone. To the contrary, ***RAID*** is used in a variety of different ways to refer to an array of disks and sometime an array of disks plus other components. As one example, the Abstract of the '346 patent defines RAID as "a redundant array of inexpensive disks," thereby referring only to the disks themselves. However, Fig. 4 of the '346 patent shows RAID 490, which includes RAID controllers 460, 471, as well as hubs 440, 441. Also, claim 9 recites that the first and second RAID controllers are "included in the RAID." In other words, even the '346 patent is inconsistent about what a ***RAID*** is.

43. Dictionary definitions tend to also be somewhat inconsistent. The Microsoft Computer Dictionary, 4^{th} ed., 1999, provides the following definition of ***RAID*** that focuses on a method:

> RAID \rad\ n. Acronym for redundant array of independent disks (formerly called redundant array of inexpensive disks). A data storage method in which data, along with information used for error correction, such as parity bits or Hamming codes, is distributed among two or more hard disks in order to improve performance and reliability. The hard disk array is governed by array management software and a disk controller, which handles the error correction. RAID is generally used on network servers. Several defined levels of RAID offer differing trade-offs among access speed, reliability, and cost. See also disk controller, error-correction coding, Hamming code, hard disk, parity bit, server (definition 1).

44. The cited art, Weygant, provides a definition of ***RAID*** that seems to focus on the disks themselves:

> RAID: RAID is an acronym for redundant array of inexpensive disks. A RAID device consists of a group of disks that can be configured in many ways, either as a single unit or in various combinations of striped and mirrored configurations. The types of configuration available are called RAID levels:
>
> • RAID 0: Disk striping.
>
> • RAID 1: Disk mirroring.
>
> • RAID 0/1: Sector Interleaved groups of mirrored disks. Also called

RAID 1/0 or RAID 10

• RAID 2: Multiple check disks using Hamming code.

• RAID 3: Byte striped, single check disk using parity.

• RAID 4: Block striped, single check disk using parity.

• RAID 5: Block striped, data and parity spread over all disks.

45. With these different definitions and uses in mind, it is my opinion that a person of ordinary skill in the art at the time would have understood the broadest reasonable interpretation of ***RAID***, in light of the present specification, to refer to "at least a redundant array of independent disks."

Challenge #1 - Claims 1-3, 5,and 8 are anticipated under 35 U.S.C. § 102(b) over Weygant

46. It is my opinion that Weygant anticipates claims 1-3, 5, and 8 of the '346 patent.

47. Weygant is a publication by Hewlett-Packard Company, describing High Availability (HA) server cluster principles in the context of Hewlett-Packard Company's UX operating system. Weygant discusses, among other things, systems having multiple server nodes where elimination of single points of failure is a goal. See, e.g., Chapter 2 of Weygant generally (discussing avoiding single points of failure). An example is the system shown in Figs. 2.10 and 2.12, where each node has redundant LAN interfaces (also referred to as "LAN cards" in Weygant), and a given LAN interface can take over for another failed LAN interface at its node. See, e.g., Weygant at p. 147, defining "LAN interface." Also, Figs. 2.10 and 2.12 illustrate that the LAN interfaces of the nodes exchange information during operation of the cluster.

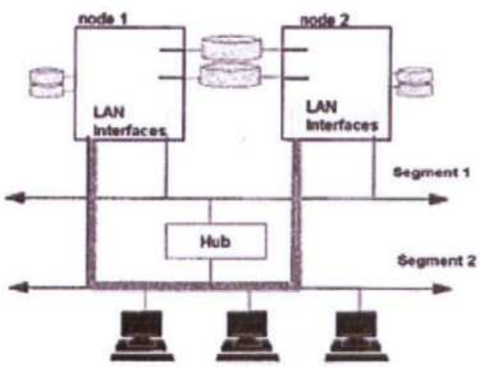

Figure 2.10 Ethernet LANs in a Grouped Subnet

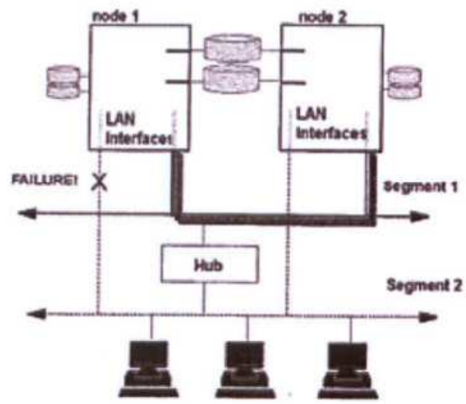

Figure 2.12 Grouped Net Following LAN Cable Failure

Weygant, Figs. 2.10 and 2.12

48. Weygant's discussion of Figs. 4.1-4.5 is very relevant because it discloses systems with redundant nodes cross-coupled with redundant hubs. The nodes of Figs. 4.1-4.5 are RAID controllers, and with that in mind, the architecture of Figs. 4.1-4.5 of Weygant is the same as that shown in Fig. 4 of the '346 patent. The redundant nodes, redundant LAN cards in the nodes, and communications between the various components of Weygant fully disclose the principles taught in the specification and claimed at claims 1-3, 5, and 8 of the '346 patent.

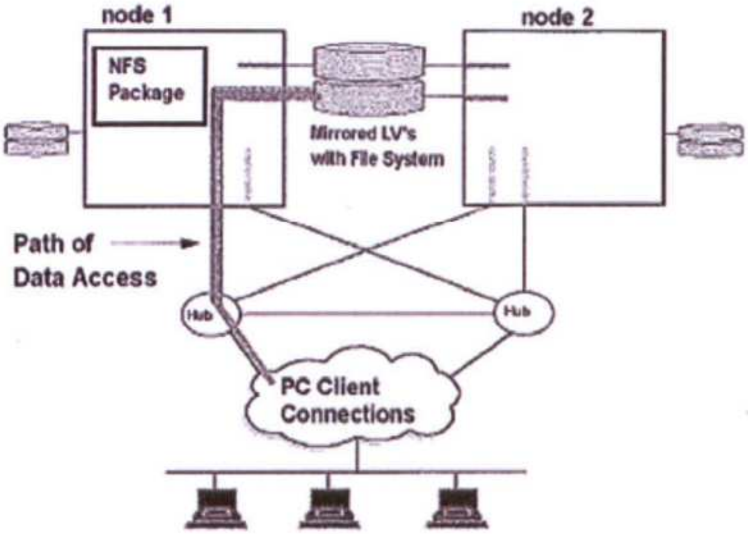

Weygant, Fig. 4.1

49. The following table explains how Weygant teaches every element of claims 1-3, 5, and 8 of the '346 patent.

50. I note here that Weygant uses the terms "LAN card" and "LAN interface" to refer to the same concepts. For instance, in the discussion at Figs. 2.10 and 2.12, Weygant discusses "LAN interfaces." By contrast, Weygant uses the term "LAN card" in Chapter 4. See Weygant at p. 115 ("an Ethernet configuration in which one LAN card on each node is active and the other is a standby."). See also Weygant at p. 147 (in the Glossary of High Availability Terminology: "**LAN interface:** The LAN interface card (LANIC) installed in a cluster node to support network services."). A person of ordinary skill in the art would have recognized that Weygant's LAN cards are LAN interfaces and that

Weygant's LAN interfaces include the LAN cards. For the purposes of this discussion, the two terms both teach a network controlling unit and a network interface controlling unit. I use both terms in the discussion below to be consistent with the particular passage to which I am referring, though Weygant makes no relevant distinction between the two terms. When I refer to a "LAN interface," my observation applies to a "LAN card" equally well (and vice versa).

Claim Language of U.S. Patent No. 6,978,346	**Relevant Disclosure in Weygant**
[1.0] An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:	Weygant discloses this limitation. See, e.g., Weygant at Figs. 4.1-4.5, showing a highly available Network File Services (NFS) system. Fig. 4.1 is reproduced below. Connections from the nodes to the hubs to the clients discloses the claimed redundant interconnection.

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant</th></tr>
<tr><td></td><td>
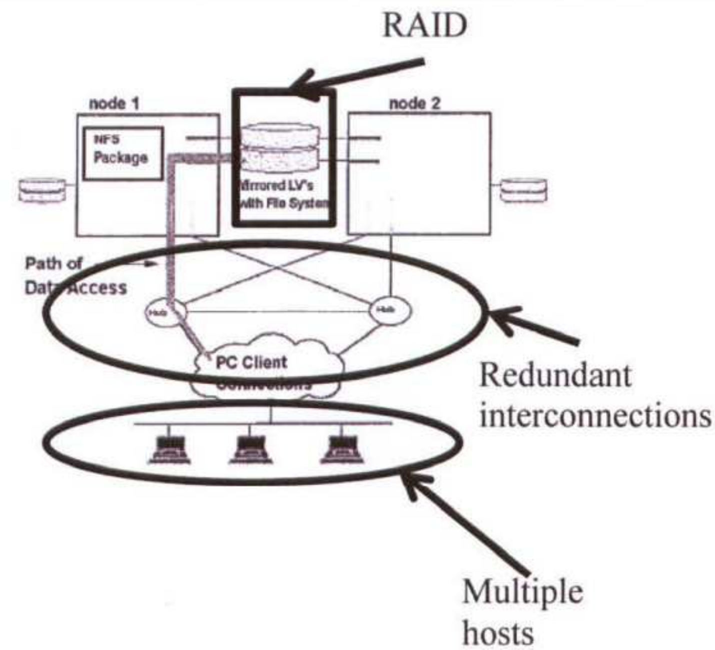

Weygant, Fig. 4.1 (annotated)

Weygant discloses that mirrored groups of disks are a RAID level 1 configuration. Weygant at 51 ("[S]oftware mirroring,[ ] is an implementation of RAID level 1 on individual disks. In HP-UX, software mirroring is created using Logical Volume Manager and the separate MirrorDisk/UX subsystem."). Figs. 4.1-4.5 of Weygant show mirrored disks. Therefore, Weygant discloses a RAID.
</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	Therefore, Weygant's system of Fig. 4.1 discloses the claimed *apparatus*; the mirrored disks disclose the claimed *RAID*, the PC clients disclose the claimed *multiple hosts*, and the connections among the hubs, nodes, and clients of Fig. 4.1 discloses the *claimed redundant interconnection*.
[1.1] a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers,	Weygant discloses this limitation. Weygant's node 1 and node 2 disclose first and second RAID controlling units, respectively. See, e.g., Weygant at various portions, disclosing that the disk mirroring is performed by software on the nodes: • p. 51 ("[A] technique for providing protected data storage is the use of software mirroring, which is an implementation of RAID level 1 on individual disks. In HP-UX, software mirroring is created using Logical Volume Manager and the separate MirrorDisk/UX subsystem."); and • p. 95 ("Basic mirroring of individual disks is provided with

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	MirrorDisk/UX. Operating through the HP-UX Logical Volume Manager"); and • p. 156 ("The use of software to provide one or more extra copies of data written to disk. This is usually done through operating system software and extensions, such as Logical Volume Manager and MirrorDisk/UX."). 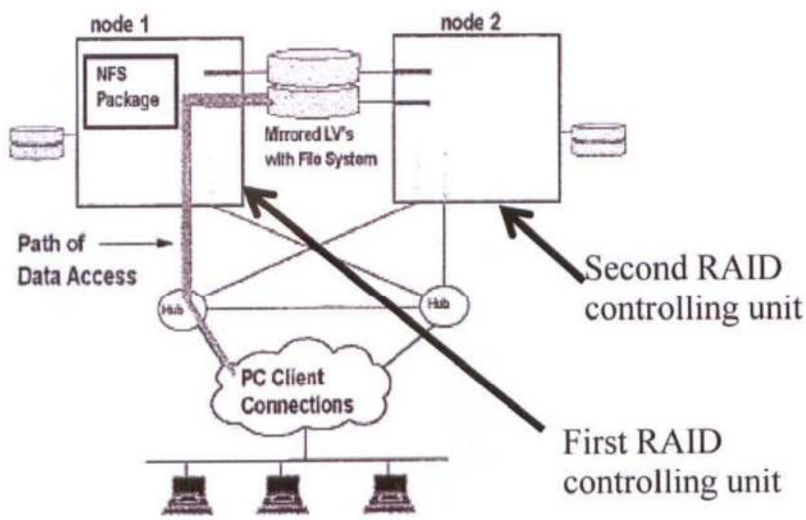 **Weygant, Fig. 4.1 (annotated)** Furthermore, in the example of Fig. 4.1, the nodes (RAID controlling units) process a requirement of the host computers by

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	providing access to a Network File Services (NFS) document system. Weygant at p. 112 ("This system uses highly available network file services (NFS). NFS is a general facility for accessing file systems remotely. In the example that follows, the NFS server software is made highly available, so that writers and editors do not lose access to their NFS mounted file systems for an extended period if the NFS server should fail. Figure 4.1 shows the basic configuration for this active/standby MC/ServiceGuard cluster."). Thus, Weygant's nodes 1 and 2 (running HP-UX Logical Volume Manager to control disk mirroring) disclose a first RAID controlling units and a second RAID controlling unit, respectively, for processing a requirement of numerous host computers.
[1.2] the first RAID controlling unit including	Weygant discloses this limitation. See, e.g., Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which <u>one LAN card on each node is active and the other is a standby</u>. The active LAN carries file server requests

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
a first network controlling unit and a second network controlling unit,	from clients and also the cluster's own heartbeat messages."), where the LAN cards disclose network controlling units. 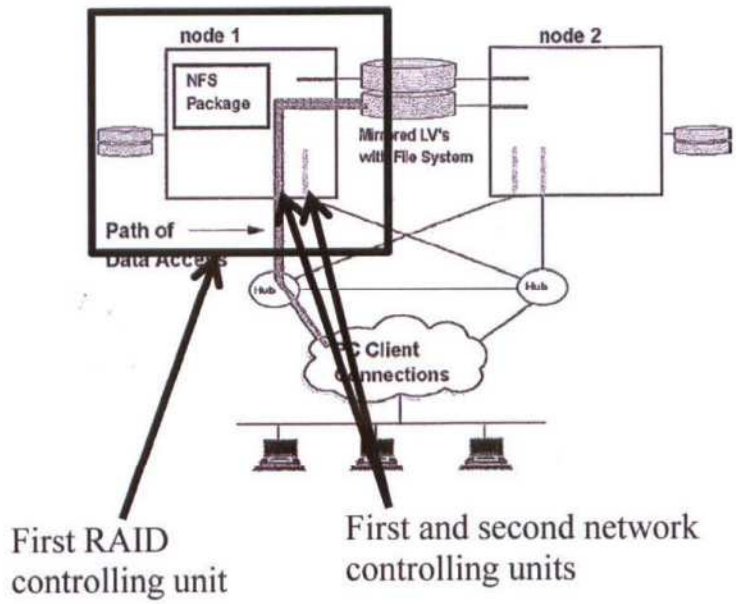 **Weygant, Fig. 4.1 (annotated)** Thus, Weygant's LAN cards in node 1 disclose a first network controlling unit and a second network controlling unit.
[1.3] and the second RAID controlling unit including	Weygant discloses this limitation. See my analysis at [1.2] (above), showing that each node (RAID controlling unit) includes two LAN cards (network controlling units). Node 2 discloses a second RAID controlling unit with

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
a third network controlling unit and a fourth network controlling unit;	third and fourth network controlling units. 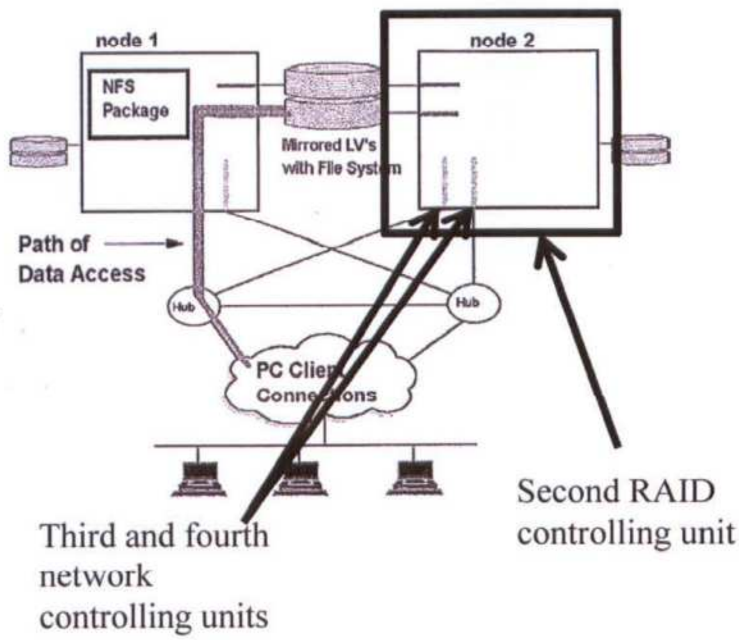 **Weygant, Fig. 4.1 (annotated)** Thus, Weygant's LAN cards in node 2 disclose a third network controlling unit and a fourth network controlling unit.
[1.4] a plurality of connection units for connecting the	Weygant discloses this limitation. See, e.g., Weygant at Fig. 4.1, showing hubs (connection units) connecting nodes 1 and 2 (first and second RAID controlling units) to the clients (host computers).

DHPN-1006

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
first RAID controlling units and the second RAID controlling unit to the numerous host computers,	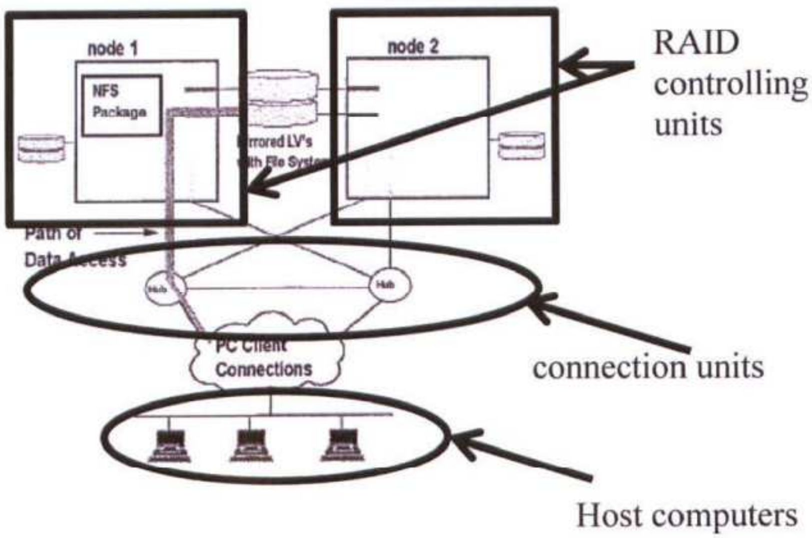 **Weygant, Fig. 4.1 (annotated)** Thus, Weygant discloses a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.
[1.5] wherein the first RAID controlling unit and the second RAID	Weygant discloses this limitation. Such information exchange is taught by Weygant at p. 118 ("Figure 4.1 shows an Ethernet configuration in which one LAN card on each node is active and the other is a standby. <u>The active LAN carries file server requests from clients</u> and also the cluster's

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
controlling unit directly exchange information with the numerous host computers through the plurality of connecting units,	own heartbeat messages."); see also Weygant at p. 112 ("the NFS server software is made highly available, so that writers and editors do not lose access to their NFS mounted file systems for an extended period if the NFS server should fail.). Also, Fig. 4.1 shows an example shaded line, illustrating information exchanged between the clients (host computers) and node 1 (first RAID controlling unit). Fig. 4.3 shows communication exchanged between clients (host computers) and node 2 (second RAID controlling unit). 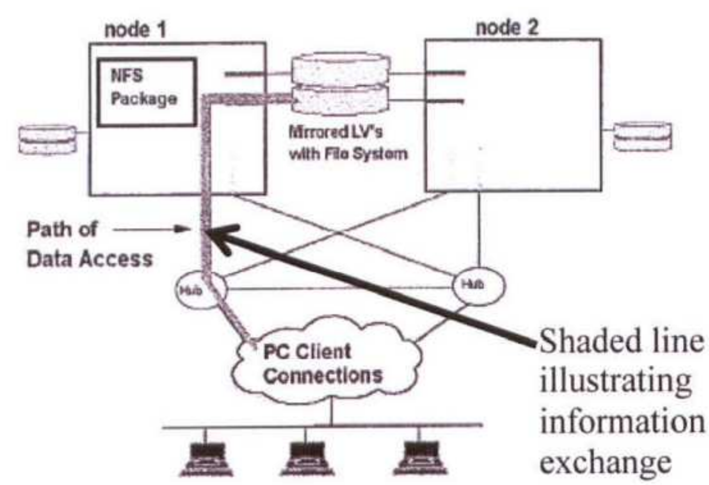

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	Weygant, Fig. 4.1 (annotated) 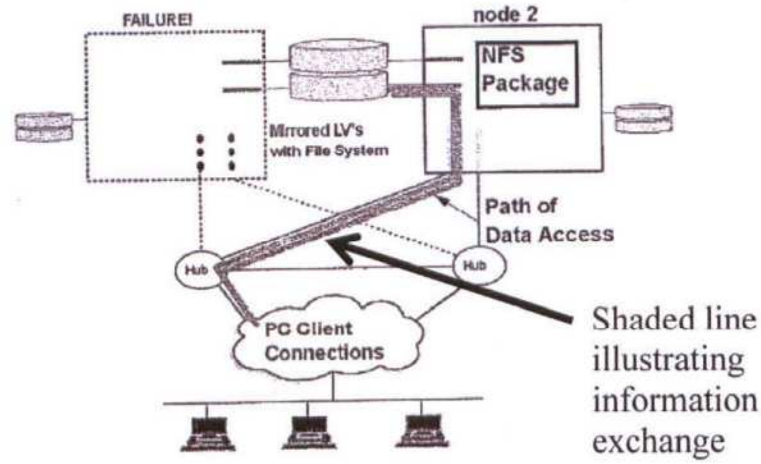 **Weygant, Fig. 4.3 (annotated)** Therefore, Weygant's nodes 1 and 2 disclose the RAID controlling units directly exchange information with the host computers through the connecting units as shown in the communication paths of Figs. 4.1-4.3.
[1.6] and the first network controlling	Weygant discloses this limitation. Weygant teaches that the nodes (RAID controllers) exchange

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
unit exchanges information with the fourth network controlling unit,	heartbeat signals via their LAN interfaces (network controlling units). • See, e.g., Weygant at p. 60 ("In a cluster, the high availability software establishes a communication link known as a heartbeat among all the nodes in the cluster on a subnet known as the heartbeat subnet. These messages allow the high availability software to tell if one or more nodes has failed."), disclosing that the nodes send heartbeat signals to each other. • See also Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which one LAN card on each node is active and the other is a standby. The active LAN carries file server requests from clients and also the cluster's own heartbeat messages."), disclosing that the heartbeats are transmitted from an active LAN interface (network controlling unit) on a node to another active LAN interface on another node. • See also Weygant at p. 123-124 ("An Ethernet

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	configuration will be used, including two LAN interfaces per node attached to different hubs.... Data and heartbeats will use one LAN interface, and an RS232 connection between the two nodes will serve as a heartbeat backup in case of heavy user traffic on the LAN. The second LAN interface will serve as a standby."), disclosing that the heartbeat signals from one node to another are transmitted by the LAN interfaces (network interface controllers). Also, Weygant teaches an active LAN interface (network interface controller) communicates with an active LAN interface of another node (RAID controller). This concept is shown in Fig. 2.10 and is applicable to the examples of Figs. 4.1-4.4.

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant</th></tr>
<tr><td></td><td>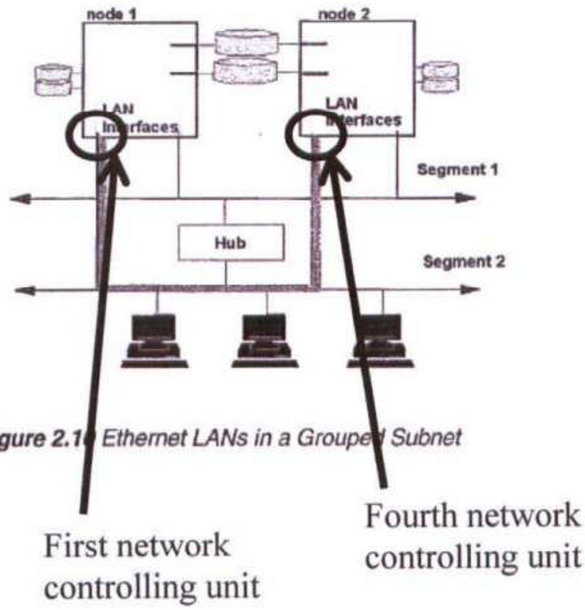

Figure 2.10 Ethernet LANs in a Grouped Subnet

Weygant, Fig. 2.10 (annotated)

In a scenario in which the first network controlling unit and the fourth network controlling unit are active, they would exchange heartbeat signals, as shown in Fig. 2.10.

Accordingly, Weygant's communication paths disclose the first network controlling unit exchanges information with the fourth network controlling unit.</td></tr>
<tr><td>[1.7] and the</td><td>Weygant discloses this feature.</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
second network controlling unit exchanges information with the third network controlling unit.	As described above at [1.6], Weygant teaches that the nodes (RAID controlling units) exchange heartbeat signals via their LAN interfaces (network controlling units). See Weygant at p. 60, 115, 123-124. Also, Weygant teaches an active LAN interface (network interface controller) communicates with an active LAN interface of another node (RAID controller). Fig. 2.12 (reproduced below) is applicable to the examples of Figs. 4.1-4.4. 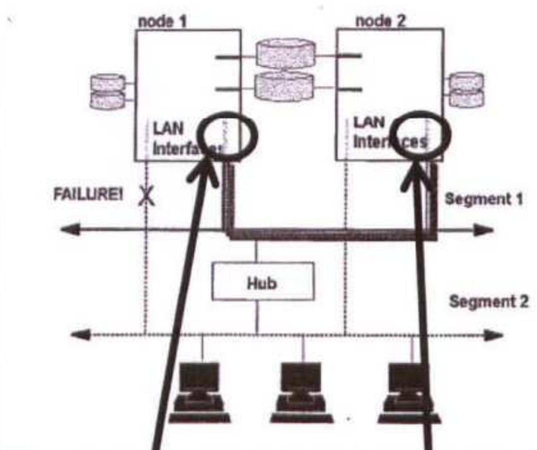 *Figure 2.12 Grouped Net Following LAN Cable Failure* Second network controlling unit Third network controlling unit

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	Weygant, Fig. 2.12 (annotated) In a scenario in which the second network controlling unit and the third network controlling unit are active, they would exchange heartbeat signals, as shown in Fig. 2.12. The designations in my annotations above (first, second, third, fourth) are exemplary, as either LAN card in node 1 can be considered first or second and either LAN card in node 2 can be considered third or fourth. Accordingly, Weygant's communication paths disclose the second network controlling unit exchanges information with the third network controlling unit.
[2.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[2.1] wherein	Weygant discloses this feature.

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant</th></tr>
<tr><td>said respective RAID controlling units are connected to the plurality of individual connecting units.</td><td>See Fig. 4.1, showing the nodes 1 and 2 (RAID controlling units) connected to the hubs (plurality of connecting units).

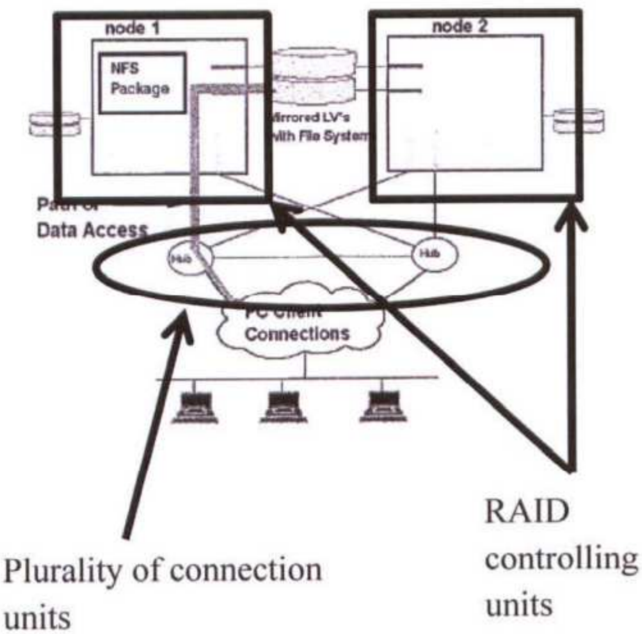

Weygant, Fig. 4.1 (annotated)

Figs. 4.2 and 4.3 (not reproduced above) show different communication paths than does Fig. 4.1, thereby illustrating various connections between the RAID controlling units and the connecting units. Thus, Weygant's nodes connected to the hubs discloses said respective RAID controlling units are connected to</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	the plurality of individual connecting units.
[3.0] The apparatus as recited in claim 2,	See my analysis at claim 2.
[3.1] wherein the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface	Weygant discloses this limitation. See, e.g., Weygant, Figure 4.1, showing LAN interfaces (first and second network controlling units) coupled to the different hubs (connection units).

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
controlling unit is coupled to the connecting unit of another side.	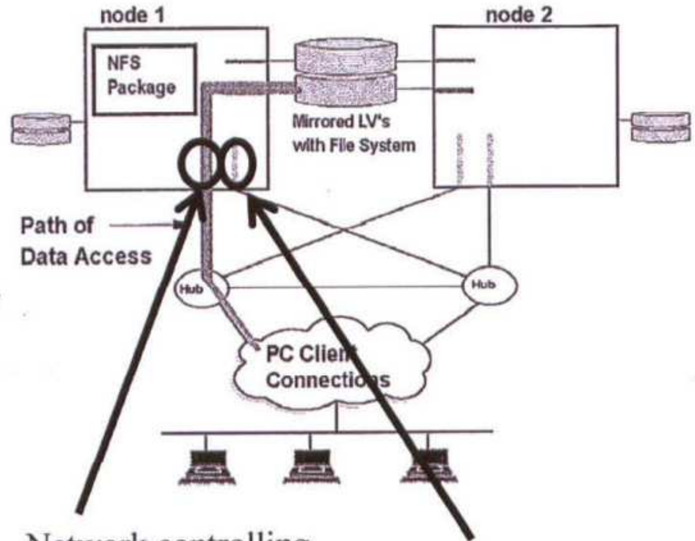 Network controlling unit coupled to connection unit (hub) on the left Network controlling unit coupled to the connection unit (hub) on the right **Weygant, Figure 4.1 (annotated)** Thus, Weygant's LAN cards respectively coupled to the hubs on the right and the left disclose the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
[5.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[5.1] wherein said plurality of connecting units have at least three connection ports,	Weygant discloses this limitation. See, e.g., Weygant, Figure 4.1, showing each hub (a connecting unit) having at least four ports—one port in communication with node 1, one port in communication with node 2, one port to the other hub, and one port to the PC client connections. Since each hub has at least four ports, the total number of ports in the plurality of connecting units is at least eight.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	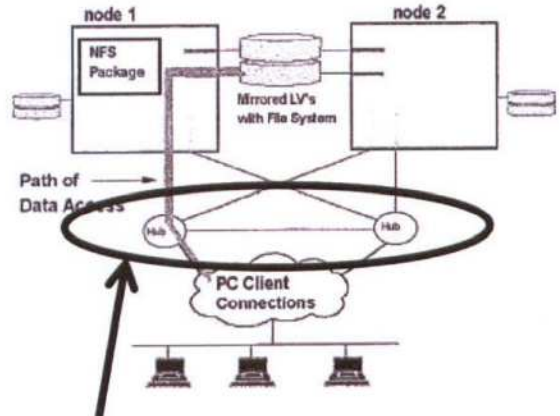 A plurality of connecting units, where each hub is shown with at least four ports **Weygant, Figure 4.1 (annotated)** Thus, Weygant's hubs with at least four ports each discloses said plurality of connecting units have at least three connection ports.
[5.2] two of the at least three connection ports is coupled to one	Weygant discloses this limitation. My analysis below demonstrates one way that the limitation may be shown. However, the designations first, second, third, and fourth in my annotations are exemplary, as either LAN card in node 1 can be considered a first or second network controlling

DHPN-1006

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant</th></tr>
<tr><td>of the first network interface controlling unit and the third network controlling unit</td><td>unit, and either LAN card in node 2 can be considered a third or a fourth network controlling unit.

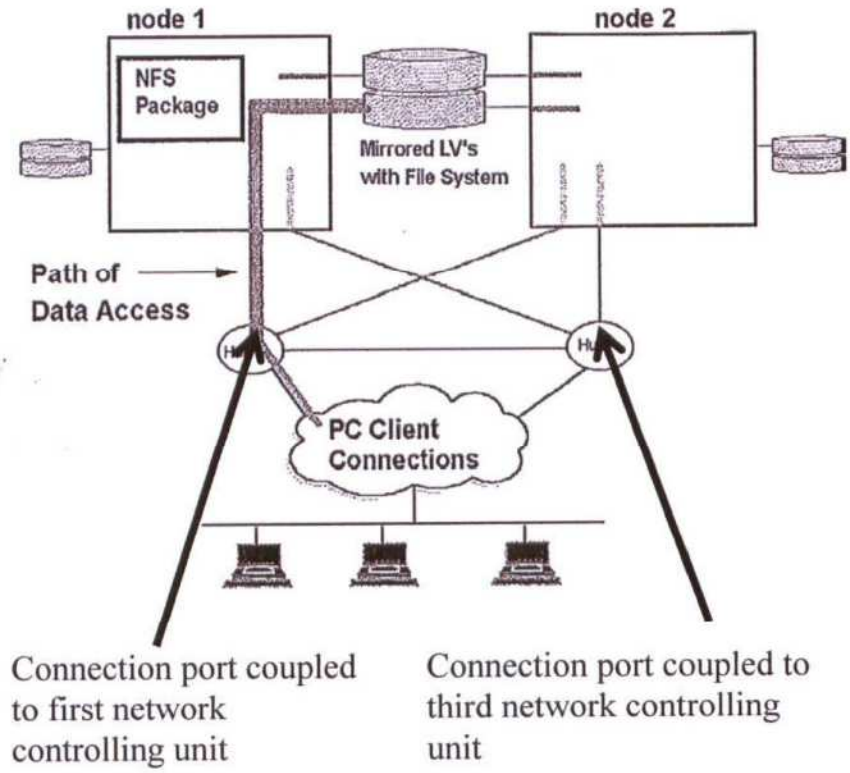

Weygant, Figure 4.1 (annotated)

Therefore, Weygant's hub ports, where one is coupled to the first LAN card and the other is coupled to the third LAN card, discloses two of the at least three connection ports is coupled to</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
	one of the first network interface controlling unit and the third network controlling unit.
[5.3] and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers.	Weygant discloses this limitation. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the clients (hosts). Weygant shows all ports in the hubs to be provided as hub equipment, therefore "the rest of the connection ports" are provided as hub equipment, and the hub equipment is connected to the host computers. Thus, the ports provided by the hubs in Weygant, where the hubs are connected with the PC clients, disclose "the rest of the connection ports being provided as a hub equipment connected with the numerous host computers."
[8.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[8.1] wherein	Weygant discloses this limitation.

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant</th></tr>
<tr><td>the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second network interface controlling unit of said first RAID</td><td>

For instance, Figure 4.1 (annotated below) shows the various connections of claim 8.

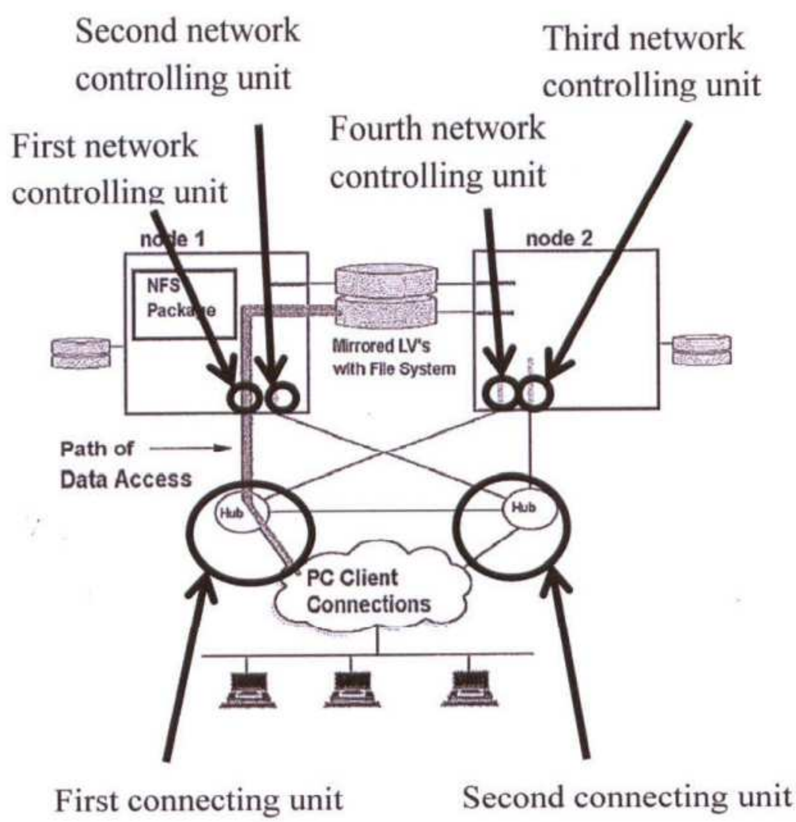

Weygant, Figure 4.1 (annotated).

The designations (first, second, third, fourth) in my annotations are exemplary, as either LAN card in node 1 can be considered a

</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
controlling unit being connected to a second connecting unit, the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit, and the fourth network	first or second network controlling unit, and either LAN card in node 2 can be considered a third or fourth network controlling unit. Thus, Weygant discloses the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit (hub on the left), the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit (hub on the right), the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit (hub on the right), and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit (hub on the left).

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant
interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.	

Challenge #2 - Claims 1-3 and 8 are obvious under 35 U.S.C. § 103(a) over Weygant in view of Mylex.

51. It is my opinion that Weygant and Mylex render obvious claims 1-3 and 8 of the '346 patent.

52. As explained above at Challenge #1, Weygant discloses nodes 1 and 2, which run software to provide disk mirroring, and disk mirroring is a RAID configuration. It my opinion that nodes 1 and 2, therefore, disclose RAID controlling units and RAID controllers. However, if someone were to argue that the nodes of Weygant are not RAID controlling units or RAID controllers, I would note that such concepts were neither new nor non-obvious in the 1999-2000 time period. It would have been obvious to a person of ordinary skill in the art that RAID controlling functionality as recited in the claims could be implemented in a variety of software and hardware configurations. For example, Mylex discloses that RAID controllers could be implemented internal or external to a host system (see Fig 1A, 1B). Mylex also discloses exemplary physical components of a RAID controller. DHPN-1007, 6:24-7:5.

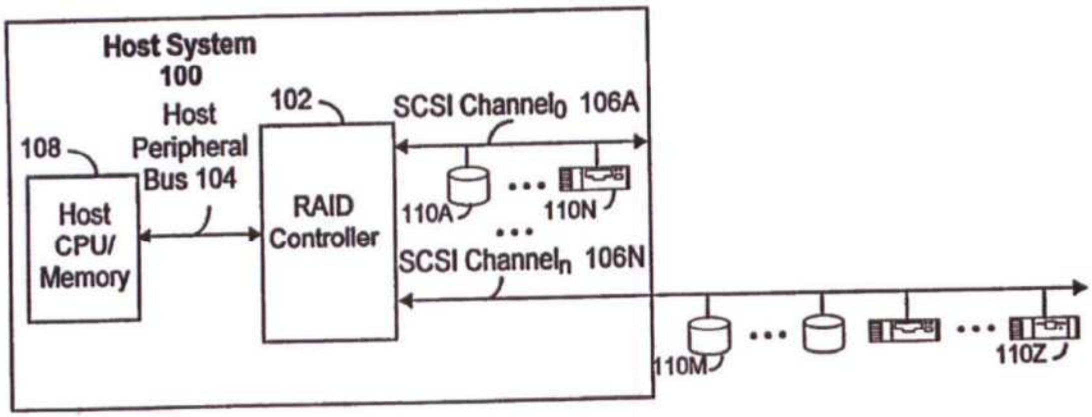

Mylex, Fig. 1A (annotated)

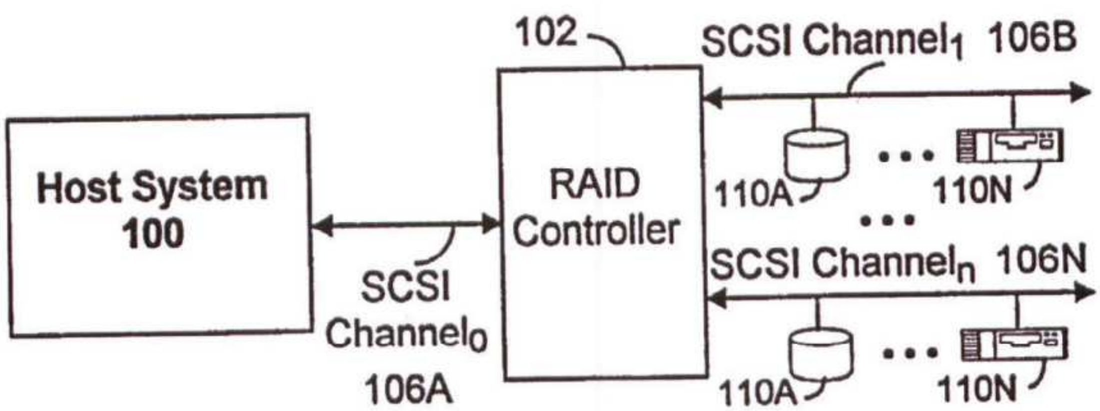

Mylex, Fig. 1B (annotated)

53. The combination of Weygant with Mylex provides the details of a component to perform RAID controlling functionality, such as the nodes disclosed in Weygant. In other words, a person of ordinary skill in the art, with the teachings of Mylex, would have understood that the RAID controlling functionality of nodes 1 and 2 could have been implemented in separate RAID controllers (or RAID controlling units), such as Mylex's controller 102, internal or external to nodes 1 and 2 of Weygant.

54. A person of ordinary skill in the art would have implemented such a combination in order to satisfy various design preferences for implementing the devices of Weygant (e.g., design preferences, such as space savings, accessibility, and cost). For instance, a person of ordinary skill in the art might have implemented a separate RAID controller, such as controller 102 of Mylex, internal or external to node 1 and/or node 2 of Weygant to take advantage of already-available RAID controller devices on the market. Further, employing any particular internal or external RAID controller configuration is merely a simple substitution of one known element for another to obtain predictable results (the predictable results including providing network connections).

55. The following table explains how Weygant and Mylex disclose every element of claims 1-3 and 8 of the '346 patent.

Claim Language of U.S. Patent No. 6,978,346	**Relevant Disclosure in Weygant and Mylex**
[1.0] An apparatus for a redundant interconnectio	Weygant discloses this limitation. See, e.g., Weygant at Figs. 4.1-4.5, showing a highly available Network File Services (NFS) system. Fig. 4.1 is reproduced below. Connections from the nodes to the hubs to the clients

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
n between multiple hosts and a RAID, comprising:	discloses the claimed redundant interconnection. 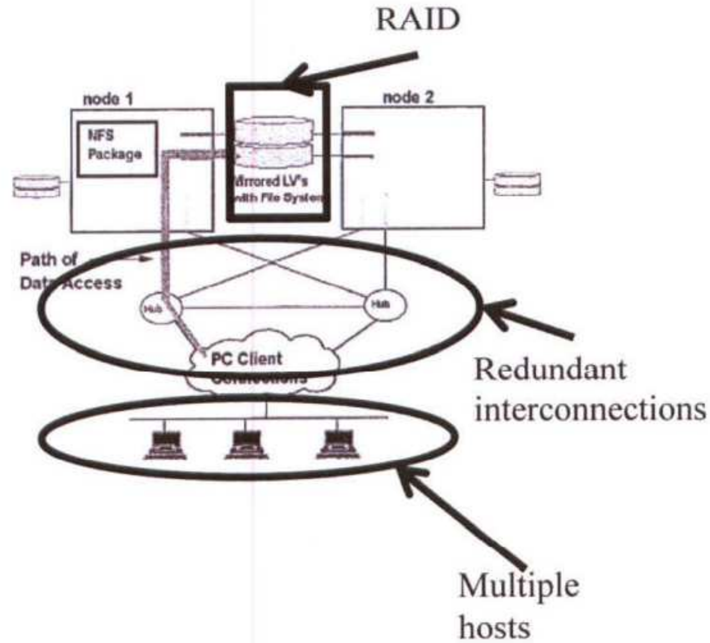 **Weygant, Fig. 4.1 (annotated)** Weygant discloses that mirrored groups of disks are a RAID level 1 configuration. Weygant at 51 ("[S]oftware mirroring,[ ] is an implementation of RAID level 1 on individual disks. In HP-UX, software mirroring is created using Logical Volume Manager and the separate MirrorDisk/UX subsystem."). Figs. 4.1-4.5 of Weygant show mirrored disks. Therefore, Weygant discloses a

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	RAID. Therefore, Weygant's system of Fig. 4.1 discloses the claimed *apparatus*; the mirrored disks disclose the claimed *RAID*, the PC clients disclose the claimed *multiple hosts*, and the connections among the hubs, nodes, and clients of Fig. 4.1 discloses the *claimed redundant interconnection.*
[1.1] a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers,	Weygant and Mylex make this limitation obvious. As I noted above at Challenge #1, Weygant's node 1 and node 2 disclose first and second RAID controlling units, respectively. See, e.g., Weygant at various portions, disclosing that the disk mirroring is performed by software on the nodes: • p. 51 ("[A] technique for providing protected data storage is the use of software mirroring, which is an implementation of RAID level 1 on individual disks. In HP-UX, software mirroring is created using Logical Volume Manager and the separate MirrorDisk/UX subsystem."); and • p. 95 ("Basic mirroring of individual disks is provided with

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	MirrorDisk/UX. Operating through the HP-UX Logical Volume Manager"); and • p. 156 ("The use of software to provide one or more extra copies of data written to disk. This is usually done through operating system software and extensions, such as Logical Volume Manager and MirrorDisk/UX."). 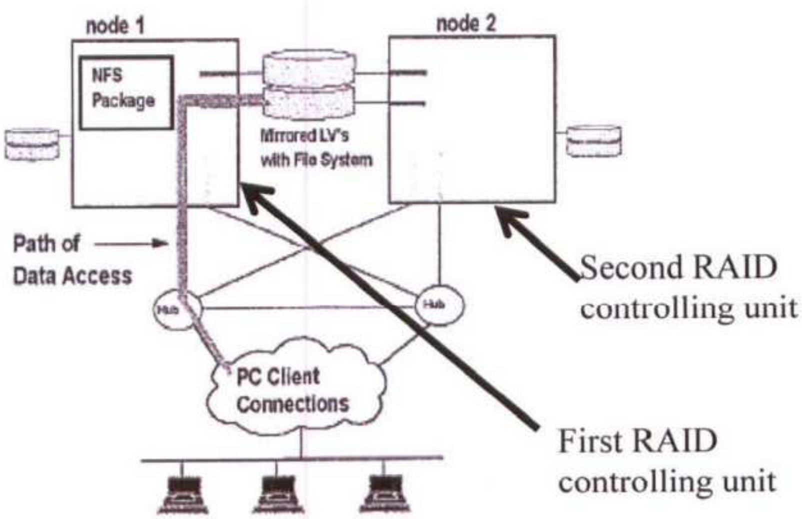 **Weygant, Fig. 4.1 (annotated)** Furthermore, in the example of Fig. 4.1, the nodes (RAID controlling units) process a requirement of the host computers by

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	providing access to a Network File Services (NFS) document system. Weygant at p. 112 ("This system uses highly available network file services (NFS). NFS is a general facility for accessing file systems remotely. In the example that follows, the NFS server software is made highly available, so that writers and editors do not lose access to their NFS mounted file systems for an extended period if the NFS server should fail. Figure 4.1 shows the basic configuration for this active/standby MC/ServiceGuard cluster."). However, if someone were to argue that nodes 1 and 2 of Weygant do not teach RAID controlling units, I would note that RAID controlling units, such as controller 102 of Mylex, were known in the art at the time. A person of ordinary skill in the art would have implemented separate RAID controllers at nodes 1 and 2 for the reasons given in the paragraphs above this chart. Thus, Weygant's nodes 1 and 2 (running HP-UX Logical Volume Manager to control disk mirroring) in view of Mylex's RAID controller disclosure teaches a first RAID controlling units and a

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	second RAID controlling unit, respectively, for processing a requirement of numerous host computers.
[1.2] the first RAID controlling unit including a first network controlling unit and a second network controlling unit,	Weygant discloses this limitation. See, e.g., Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which one LAN card on each node is active and the other is a standby. The active LAN carries file server requests from clients and also the cluster's own heartbeat messages."), where the LAN cards disclose network controlling units. 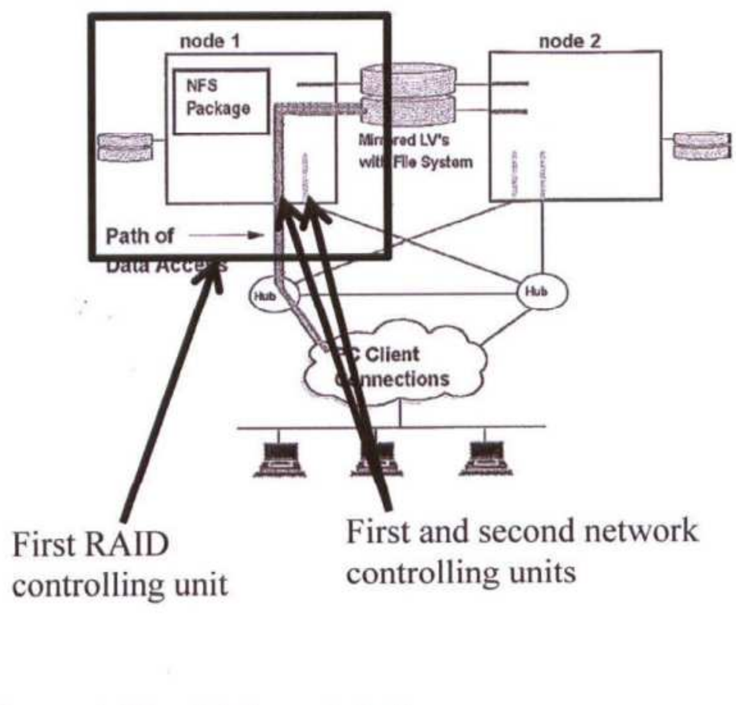 **Weygant, Fig. 4.1 (annotated)**

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	Thus, Weygant's LAN cards in node 1 disclose a first network controlling unit and a second network controlling unit.
[1.3] and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit;	Weygant discloses this limitation. See, analysis at [1.2] (above), showing that each node (RAID controlling unit) includes two LAN cards (network controlling units). Node 2 discloses a second RAID controlling unit with third and fourth network controlling units. 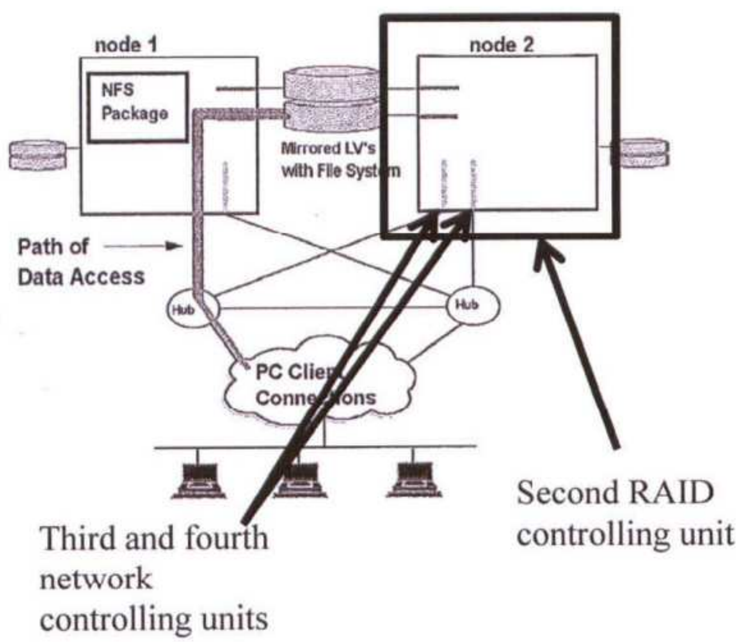 **Weygant, Fig. 4.1 (annotated)** Thus, Weygant's LAN cards in node 2 disclose a third network

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	controlling unit and a fourth network controlling unit.
[1.4] a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers,	Weygant discloses this limitation. See, e.g., Weygant at Fig. 4.1, showing hubs (connection units) connecting nodes 1 and 2 (first and second RAID controlling units) to the clients (host computers). 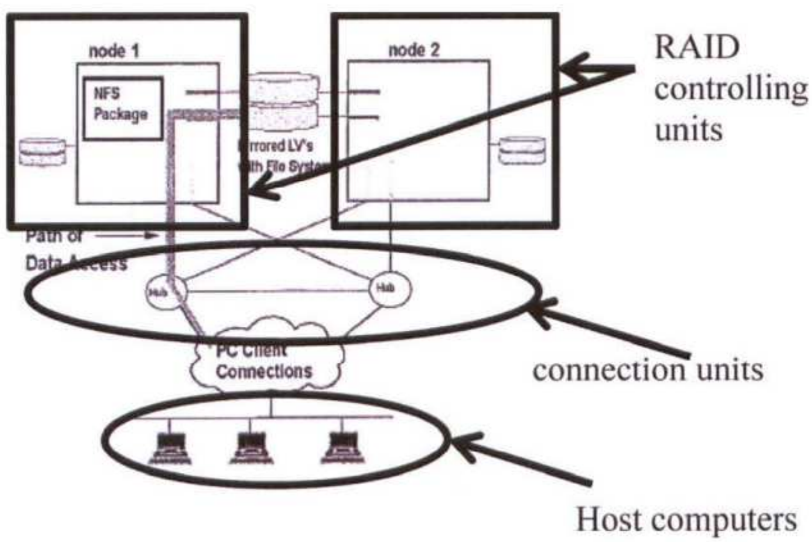 **Weygant, Fig. 4.1 (annotated)** Thus, Weygant discloses a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
[1.5] wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units,	Weygant discloses this limitation. Such information exchange is taught by Weygant at p. 118 ("Figure 4.1 shows an Ethernet configuration in which one LAN card on each node is active and the other is a standby. The active LAN carries file server requests from clients and also the cluster's own heartbeat messages."); see also Weygant at p. 112 ("the NFS server software is made highly available, so that writers and editors do not lose access to their NFS mounted file systems for an extended period if the NFS server should fail."). Also, Fig. 4.1 shows an example shaded line, illustrating information exchanged between the clients (host computers) and node 1 (first RAID controlling unit). Fig. 4.3 shows communication exchanged between clients (host computers) and node 2 (second RAID controlling unit).

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	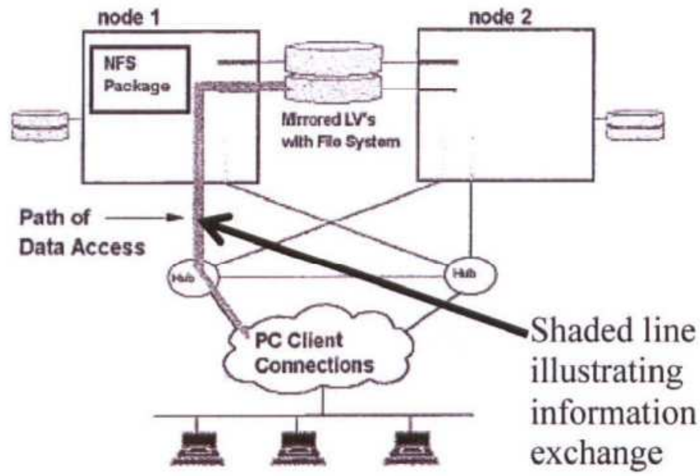 **Weygant, Fig. 4.1 (annotated)** 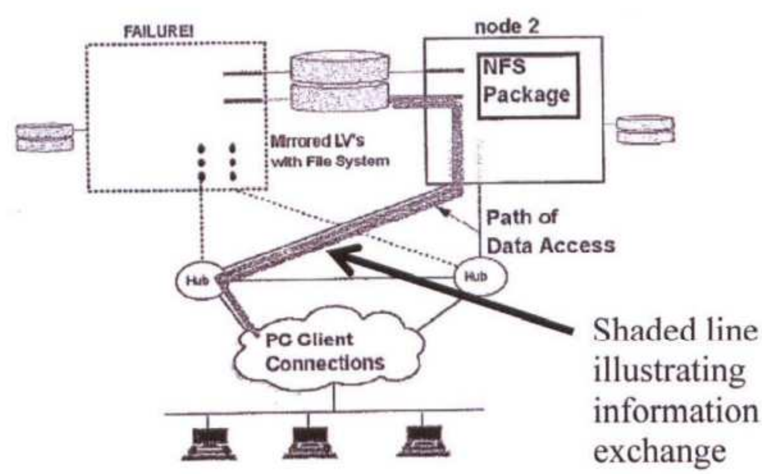

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	Weygant, Fig. 4.3 (annotated) Therefore, Weygant's nodes 1 and 2 disclose the RAID controlling units directly exchange information with the host computers through the connecting units as shown in the communication paths of Figs. 4.1-4.3.
[1.6] and the first network controlling unit exchanges information with the fourth network controlling unit,	Weygant discloses this limitation. Weygant teaches that the nodes (RAID controllers) exchange heartbeat signals via their LAN interfaces (network controlling units). • See, e.g., Weygant at p. 60 ("In a cluster, the high availability software establishes <u>a communication link known as a heartbeat among all the nodes in the cluster</u> on a subnet known as the heartbeat subnet. These messages allow the high availability software to tell if one or more nodes has failed."), disclosing that the nodes send heartbeat signals to each other. • See also Weygant at p. 115 ("Figure 4.1 shows an Ethernet

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	configuration in which one LAN card on each node is active and the other is a standby. The active LAN carries file server requests from clients and also the cluster's own heartbeat messages."), disclosing that the heartbeats are transmitted from an active LAN interface (network controlling unit) on a node to another active LAN interface on another node. • See also Weygant at p. 123-124 ("An Ethernet configuration will be used, including two LAN interfaces per node attached to different hubs.... Data and heartbeats will use one LAN interface, and an RS232 connection between the two nodes will serve as a heartbeat backup in case of heavy user traffic on the LAN. The second LAN interface will serve as a standby."), disclosing that the heartbeat signals from one node to another are transmitted by the LAN interfaces (network interface controllers). Also, Weygant teaches an active LAN interface (network interface controller) communicates with an active LAN interface

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	of another node (RAID controller). This concept is shown in Fig. 2.10 and is applicable to the examples of Figs. 4.1-4.4. 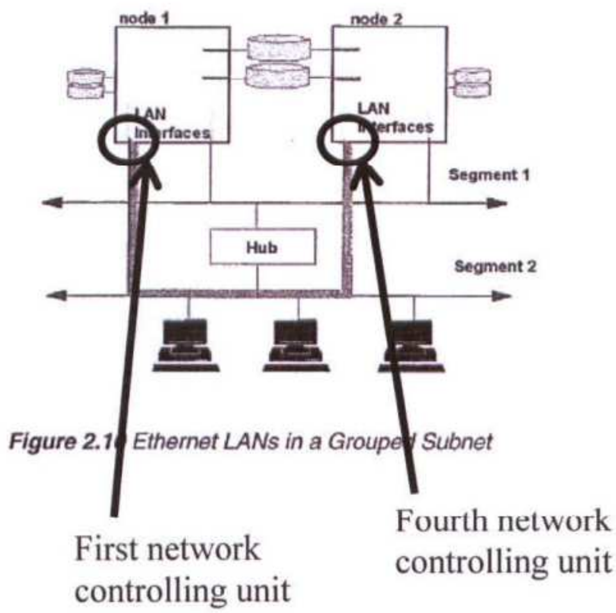 *Figure 2.10 Ethernet LANs in a Grouped Subnet* **Weygant, Fig. 2.10 (annotated)** In a scenario in which the first network controlling unit and the fourth network controlling unit are active, they would exchange heartbeat signals, as shown in Fig. 2.10. Accordingly, Weygant's communication paths disclose the first

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	network controlling unit exchanges information with the fourth network controlling unit.
[1.7] and the second network controlling unit exchanges information with the third network controlling unit.	Weygant discloses this feature. As described above at [1.6], Weygant teaches that the nodes (RAID controlling units) exchange heartbeat signals via their LAN interfaces (network controlling units). See Weygant at p. 60, 115, 123-124. Also, Weygant teaches an active LAN interface (network interface controller) communicates with an active LAN interface of another node (RAID controller). Fig. 2.12 (reproduced below) is applicable to the examples of Figs. 4.1-4.4.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	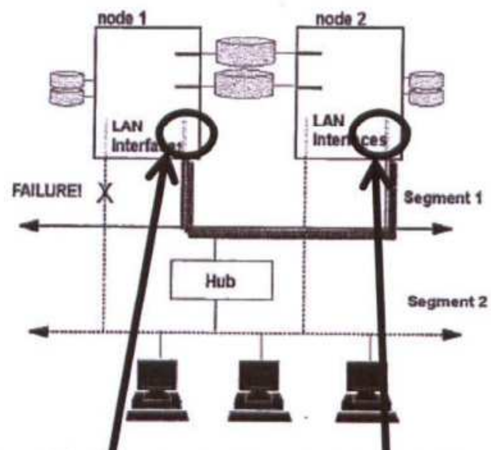 *Figure 2.12 Grouped Net Following LAN Cable Failure* Second network controlling unit Third network controlling unit **Weygant, Fig. 2.12 (annotated)** In a scenario in which the second network controlling unit and the third network controlling unit are active, they would exchange heartbeat signals, as shown in Fig. 2.12. The designations in my annotations above (first, second, third, fourth) are exemplary, as either LAN card in node 1 can be considered first or second and any LAN card in node 2 can be considered third or fourth.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	Accordingly, Weygant's communication paths disclose the second network controlling unit exchanges information with the third network controlling unit.
[2.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[2.1] wherein said respective RAID controlling units are connected to the plurality of individual connecting units.	Weygant discloses this feature. See Fig. 4.1, showing the nodes 1 and 2 (RAID controlling units) connected to the hubs (plurality of connecting units).

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
	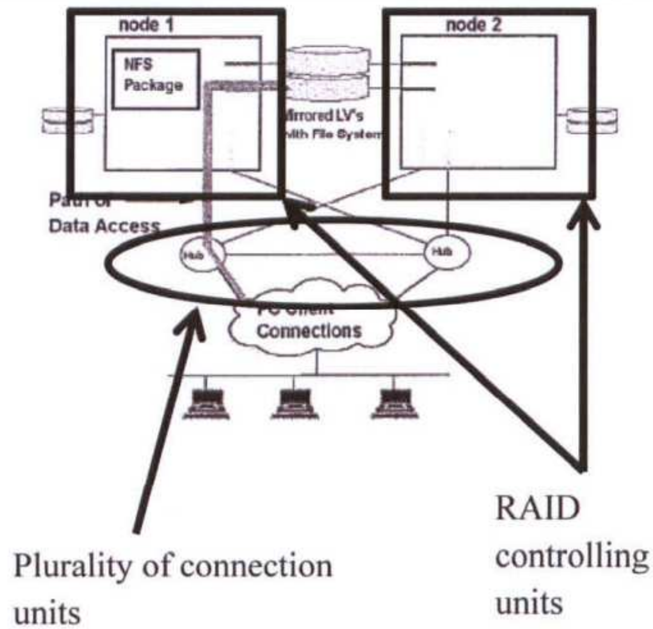 **Weygant, Fig. 4.1 (annotated)** Figs. 4.2 and 4.3 (not reproduced above) show different communication paths than does Fig. 4.1, thereby illustrating various connections between the RAID controlling units and the connecting units. Thus, Weygant's nodes connected to the hubs discloses said respective RAID controlling units are connected to the plurality of individual connecting units.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
[3.0] The apparatus as recited in claim 2,	See my analysis at claim 2.
[3.1] wherein the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling	Weygant discloses this limitation. See, e.g., Weygant, Figure 4.1, showing LAN interfaces (first and second network controlling units) coupled to the different hubs (connection units).

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
unit is coupled to the connecting unit of another side.	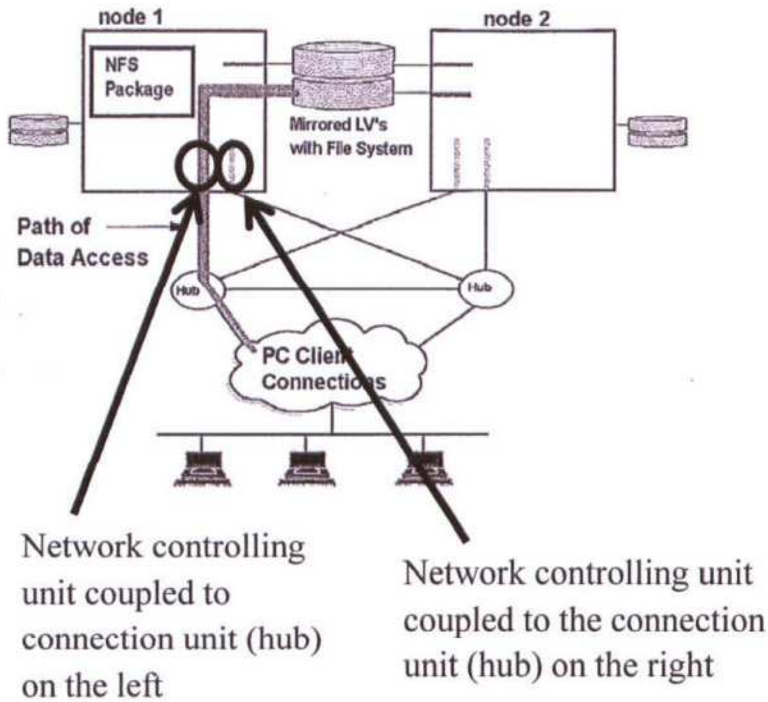 **Weygant, Figure 4.1 (annotated)** Thus, Weygant's LAN cards respectively coupled to the hubs on the right and the left disclose the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
[8.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[8.1] wherein the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second	Weygant discloses this limitation. For instance, Figure 4.1 (annotated below) shows the various connections of claim 8.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit, the third network interface controlling unit of the second RAID controlling unit being	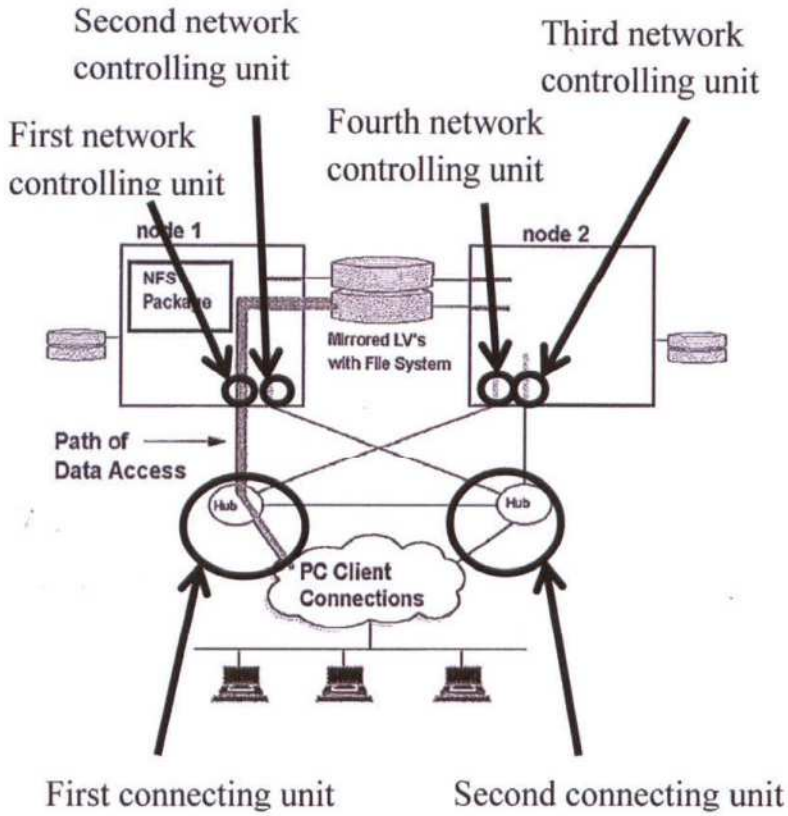 **Weygant, Figure 4.1 (annotated).** The designations (first, second, third, fourth) in my annotations are exemplary, as either LAN card in node 1 can be considered a first or second network controlling unit, and either LAN card in node 2 can be considered a third or fourth network controlling unit.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant and Mylex
connected to the second connecting unit, and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.	Thus, Weygant discloses the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit (hub on the left), the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit (hub on the right), the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit (hub on the right), and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit (hub on the left).

Challenge #3: Claims 4 and 9 are obvious under 35 U.S.C. § 103(a) over

Weygant and Mylex in view of ServiceGuard

56. It is my opinion that Weygant in view of Mylex, when modified as proposed using concepts from ServiceGuard, renders obvious claims 4 and 9 of the '346 patent.

57. ServiceGuard is another Hewlett-Packard Company publication. ServiceGuard describes the MC/ServiceGuard software, available from Hewlett-Packard Company, to be used with the HP-UX operating system on nodes of HA clusters. See, e.g., ServiceGuard at p. 16, providing an overview of the ServiceGuard software.

58. ServiceGuard discloses that active and standby network interface controllers are determined at configuration. See ServiceGuard at p. 70-71 ("It is recommended that you configure heartbeats on all subnets, including those to be used for client data. On the worksheet, enter the following for each LAN interface....IP Address: Enter this node's host IP address intended to be used on this interface. The IP address is a string of digits separated with periods in the form 'nnn.nnn.nnn.nnn'. If the interface is a standby and does not have an IP address, enter 'Standby.'). Thus, one principle taught by ServiceGuard is that a given LAN interface in a node can be a standby interface or an active interface as a matter of system configuration.

59. One skilled in the art would have multiple reasons to combine the teachings of Weygant with ServiceGuard. For instance, ServiceGuard specifically states that its software package is to be used with the MirrorDisk/UX software of Weygant. See, e.g., ServiceGuard at 18 ("MC/ServiceGuard is designed to work in conjunction with other Hewlett-Packard Company high availability products, such as MirrorDisk/UX, which provides disk redundancy to eliminate single points of failure in the disk subsystem....). Thus, ServiceGuard itself provides an explicit suggestion to combine.

60. Weygant teaches using LAN interfaces (a type of network controlling unit) in high-availability computing nodes to communicate among the nodes. As shown above in more detail for Challenge #2 at [1.1], the nodes themselves are used as RAID controllers (or RAID controlling units), and Mylex teaches that the RAID controlling functionality of nodes 1 and 2 can be embodied as separate RAID controlling units internal or external to the nodes. Weygant teaches that at each node, a LAN interface is active and another LAN interface is a standby interface. Weygant at p. 115. It is proposed to adopt from ServiceGuard the concept that a given LAN interface in a node can be configured to be active or standby. In the proposed combination, any given LAN interface in one node can exchange information (e.g., heartbeat signals) with an active LAN interface in another node or communicate with a host computer; it just depends on which LAN interfaces are configured to be active. Therefore, for any given subset of LAN interfaces, the designation of various ones as active or standby would have been merely a configuration decision within the abilities of a person of ordinary skill in the art in 2000.

61. Thus, for a given subset of LAN interfaces (such as those in Weygant), ServiceGuard teaches different configurations for the LAN interface. There is no functional distinction between which given LAN interface is configured to provide a particular functionality. Implementing this combination would have been a simple configuration decision within the abilities of a person of ordinary skill in the art in 2000. Also, it would have been a simple configuration decision that would have enabled improved redundancy and fault tolerance in a given system.

62. Furthermore, it would have been obvious to combine the teachings of Weygant with the teachings of ServiceGuard because it is a combination of known elements that achieve the predictable results of network controlling units exchanging information with each other over an available pathway.

63. The following table explains how the modification of Weygant and Mylex using the concepts of ServiceGuard teaches every element of claims 4 and 9 of the '346 patent.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
[4.0] The apparatus as recited in claim 3, wherein	See my analysis at claim 3 in the Challenge #2 invalidity chart.
[4.1] the first network interface controlling unit and the third network interface controlling unit process the requirement of the numerous	Weygant in view of ServiceGuard makes this feature obvious. See, e.g., Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which <u>one LAN card on each node is active and the other is a standby. The active LAN carries file server requests from clients</u> and also the cluster's own heartbeat messages."), where the LAN cards disclose network controlling units. Such passage also teaches that the network controlling units process requests from the clients (host computers) by carrying file server requests. As explained above, whether the first or second LAN card (network controlling unit) is active at node 1 (first RAID controlling unit) is a matter of configuration, as evidenced by

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
host computers; and	ServiceGuard. Similarly, whether the third or fourth LAN card (network controlling unit) is active at node 2 (second RAID controlling unit) is a matter of configuration, as evidenced by ServiceGuard. Therefore, in a scenario in which the first and third LAN cards (network controlling units) are configured to be active, the first and third LAN cards process requirements of the clients (host computers). Thus, Weygant in view of ServiceGuard discloses the first network interface controlling unit and the third network interface controlling unit process the requirement of the numerous host computers.
[4.2] the second network interface controlling unit and the	Weygant in view of ServiceGuard makes this feature obvious. See, e.g., Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which <u>one LAN card on each node is active and the other is a standby. The active LAN carries</u> file server requests from clients and also<u> the cluster's own heartbeat messages</u>."), where the LAN cards disclose network controlling units. Such

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
fourth network controlling unit are used for communication between the first RAID controlling unit and the second RAID controlling unit when the first and second RAID controlling units are not faulty and	passage also teaches that the network controlling units communicate between node 1 and node 2 (first and second RAID control units) by passing heartbeat signals therebetween. As explained above, whether the first or second LAN card (network controlling unit) is active at node 1 (first RAID controlling unit) is a matter of configuration, as evidenced by ServiceGuard. Similarly, whether the third or fourth LAN card (network controlling unit) is active at node 2 (second RAID controlling unit) is a matter of configuration, as evidenced by ServiceGuard. Therefore, in a scenario in which the second and fourth LAN cards (network controlling units) are configured to be active, the second and fourth LAN cards would pass the heartbeat signals. Thus, Weygant in view of ServiceGuard discloses the second network interface controlling unit and the fourth network controlling unit are used for communication between the first RAID controlling unit and the second RAID controlling unit

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	when the first and second RAID controlling units are not faulty.
[4.3] the second network interface controlling unit and the fourth network controlling unit are used for executing a function of the first network interface controlling unit and the third network controlling	Weygant discloses this feature by disclosing a failover operation when one node is faulty. In Weygant, when one node is faulty, the surviving node performs the functions of the faulty node. See, e.g., Weygant at p. 119 ("If the SPU experiences a failure, or if the operating system experiences a panic (fatal error), the node will shut down, and MC/ServiceGuard on the other node will start the package in its alternate location. The failover should take about 45 seconds in addition to the time required to start NFS.") See also Weygant at p. 125 ("In the event of SPU failure, the applications will continue running on the alternate node until the appropriate repair can be made on the failed node. After the loss of a node, of course, services will not be highly available until the repaired node re-enters the cluster."). Therefore, Weygant discloses that when a node (e.g., the first RAID controlling unit) fails, the other node (the second RAID controlling unit) runs the application package of the failed node.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty.	Cluster failover is illustrated in Weygant at Figs. 3.1 (before failover) and 3.2 (after failover), the principles of which are applicable to the examples of Figs. 4.1-4.4. 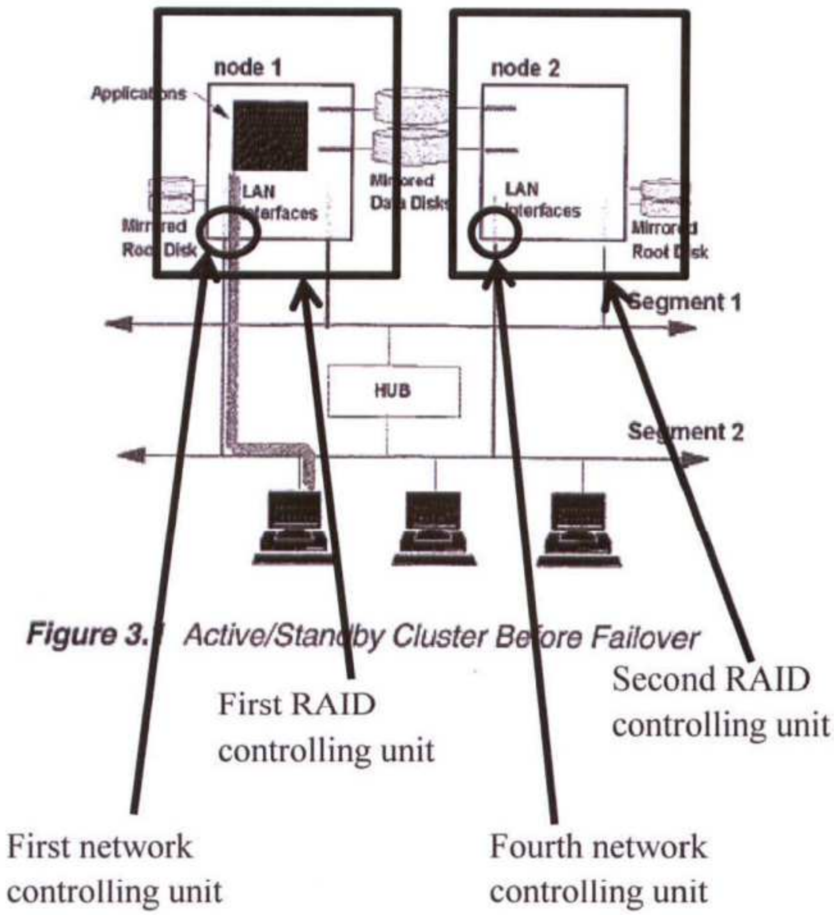 **Weygant, Fig. 3.1 (annotated)**

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	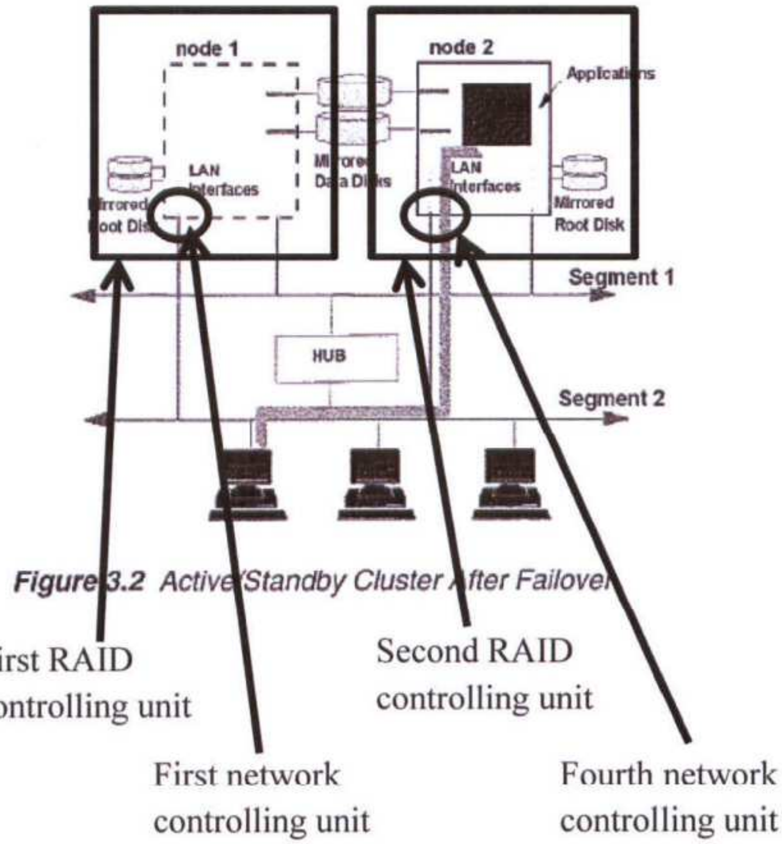 **Weygant, Fig. 3.2 (annotated)** As illustrated by Figs. 3.1 and 3.2, in a scenario in which the first network controlling unit is active and in which the fourth network controlling unit is active, a failure of the first RAID controlling unit will result in the second RAID controlling unit (and its active

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	network controlling unit—the fourth) taking over. Thus, Weygant in view of ServiceGuard discloses the fourth network interface controller in the second RAID controller is used for fault tolerance by performing functions of the first network interface controller in the first RAID controller when the first RAID controller is faulty. Of course, the concepts of Figs. 3.1 and 3.2 apply to the embodiments of Wygant's Figs. 4.1-4.5 to each node and each LAN card, regardless of which node is designated "first" or "second" or which LAN card is designated "first," "second," "third," or "fourth." Therefore, the concepts of Fig. 3.1 and 3.2 of Weygant also show that the second network interface controller in the first RAID controller is used for fault tolerance by performing functions of the third network interface controller in the second RAID controller when the second RAID controller is faulty.
[9.0] An apparatus for a	Weygant discloses this feature. See, e.g., Weygant at Figs. 4.1-4.5, showing a highly available

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
redundant interconnection between multiple host computers and a RAID, the apparatus comprising:	Network File Services (NFS) system. Fig. 4.1 is reproduced below. Connections from the nodes to the hubs to the clients disclose the claimed redundant interconnection. 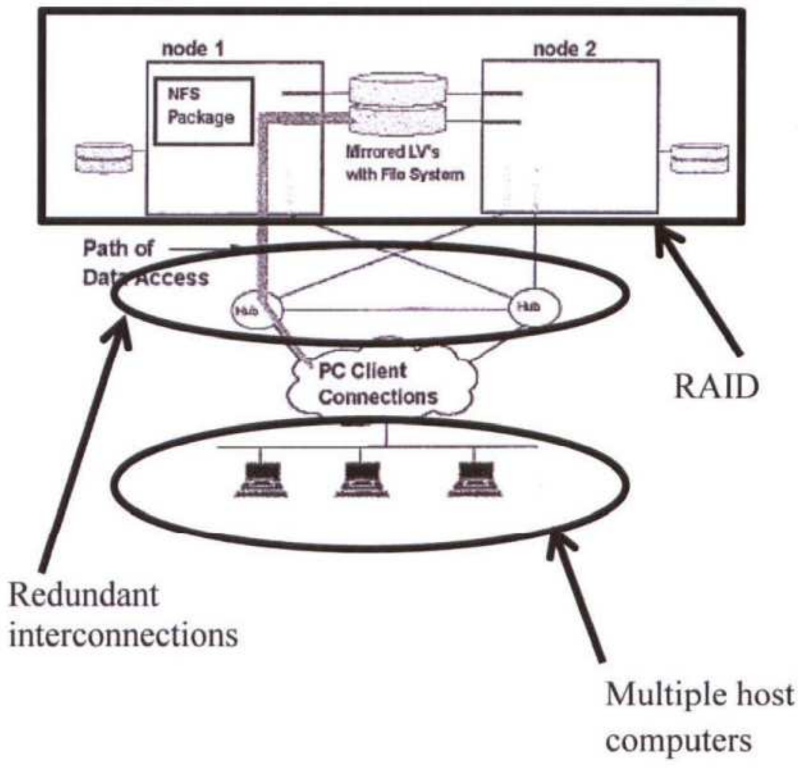 **Weygant, Fig. 4.1 (annotated)** Weygant discloses a RAID, as annotated above in Fig. 4.1. Weygant discloses that mirrored groups of disks are a RAID level

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	1 configuration. Weygant at p. 51 ("[S]oftware mirroring,[ ] is an implementation of RAID level 1 on individual disks. In HP-UX, software mirroring is created using Logical Volume Manager and the separate MirrorDisk/UX subsystem."). Therefore, Weygant discloses a RAID. Weygant's PC clients disclose multiple hosts. Thus, Weygant discloses an apparatus for a redundant interconnection between multiple host computers and a RAID.
[9.1] a plurality of connection units for connecting the host computers and the RAID;	Weygant discloses this feature. Weygant at Fig. 4.1 shows two hubs, where each hub discloses a connection unit. As noted above, the nodes and the mirrored sets of disks, collectively, disclose the RAID.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	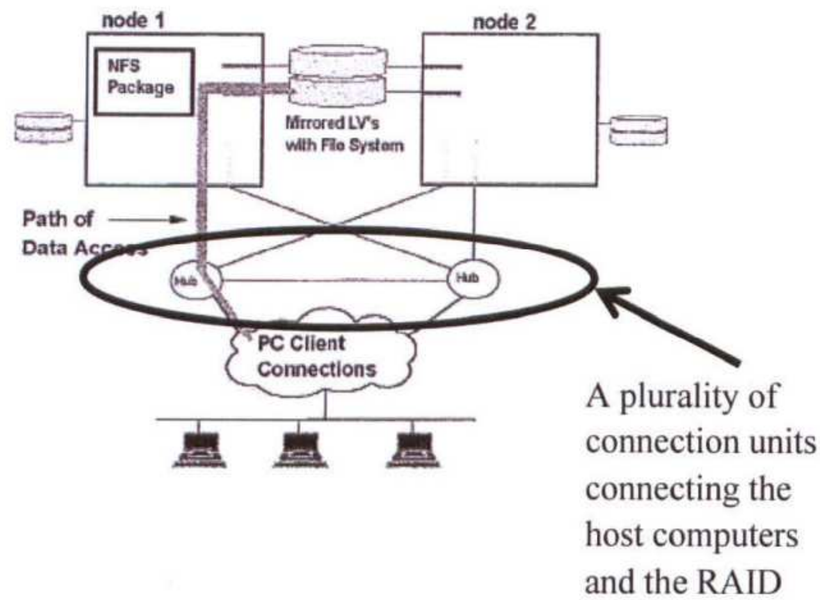 **Weygant, Fig. 4.1 (annotated)** Thus, Weygant discloses a plurality of connection units for connecting the host computers and the RAID.
[9.2] a first and a second RAID controllers, included in the	Weygant and Mylex render obvious this feature. Weygant's node 1 and node 2 disclose first and second RAID controllers, respectively. Weygant, at various portions, discloses that the disk mirroring is performed by software on the nodes (in

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
RAID,	other words, the nodes are providing control of the RAID): • p. 51 ("[A] technique for providing protected data storage is the use of software mirroring, which is an implementation of RAID level 1 on individual disks. In HP-UX, software mirroring is created using Logical Volume Manager and the separate MirrorDisk/UX subsystem."); • p. 95 ("Basic mirroring of individual disks is provided with MirrorDisk/UX. Operating through the HP-UX Logical Volume Manager"); and • p. 156 ("The use of software to provide one or more extra copies of data written to disk. This is usually done through operating system software and extensions, such as Logical Volume Manager and MirrorDisk/UX.").

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant, Mylex, and ServiceGuard</th></tr>
<tr><td></td><td>
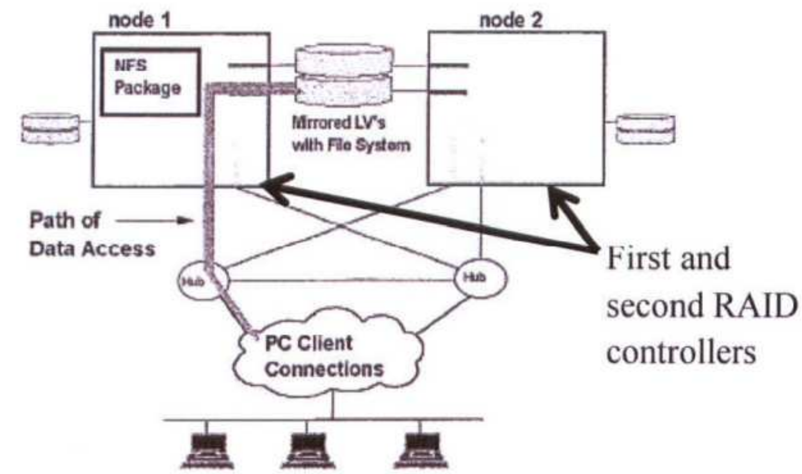

Weygant, Fig. 4.1 (annotated)

Therefore, Weygant's node 1 and node 2 (running HP-UX Logical Volume Manager to control disk mirroring) disclose first and second RAID controllers, respectively. However, if someone were to argue that nodes 1 and 2 of Weygant do not teach RAID controlling units, I would note that RAID controlling units, such as controller 102 of Mylex, were known in the art at the time. A person of ordinary skill in the art could have implemented separate RAID controllers at nodes 1 and 2 for the reasons given in the paragraphs at Challenge #2 (e.g., to satisfy various design
</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	preferences and because it is a simple substitution of one element for another to obtain predictable results). Thus, Weygant and Mylex disclose a first and a second RAID controllers, included in the RAID.
[9.3] each of which having a first network interface controller and a second network interface controller for processing requests from the plurality of the host	Weygant discloses this feature. See, e.g., Weygant at 115 ("Figure 4.1 shows an Ethernet configuration in which one LAN card on each node is active and the other is a standby. The active LAN carries file server requests from clients and also the cluster's own heartbeat messages."), where the two LAN cards disclose network interface controllers. This passage also teaches that the network interface controllers process requests from the clients (plurality of host computers) by carrying file server requests. Also, Fig. 4.1 shows the LAN cards connected to the PC clients by the hubs (connection units).

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
computers connected through the plurality of the connection units,	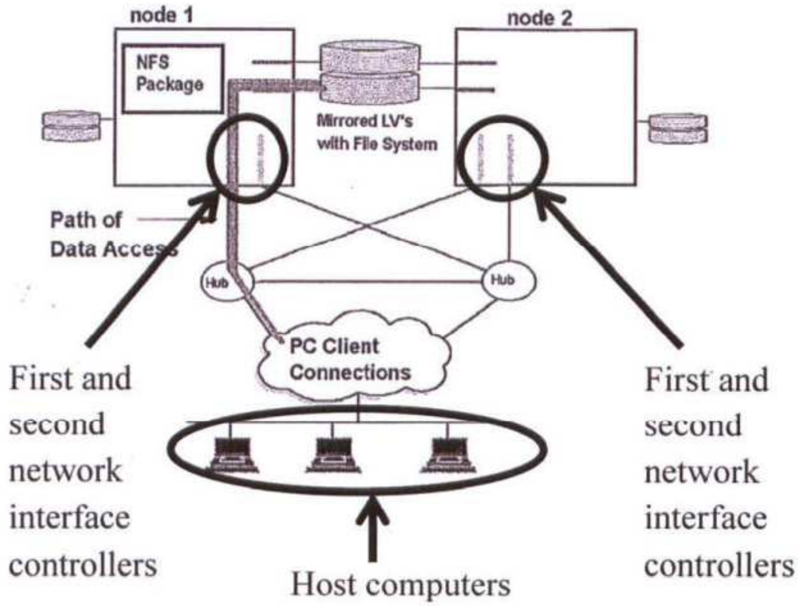 **Weygant, Fig. 4.1 (annotated)** Thus, Weygant discloses [the first and second RAID controllers] having a first network interface controller and a second network interface controller for processing requests from the plurality of the host computers connected through the plurality of the connection units.
[9.4] wherein the first	Weygant discloses this feature.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
network interface controller in the first RAID controller supplies data to the host computers connected through the plurality of connection units and	Weygant at Fig. 4.1 (reproduced and annotated below) shows a first RAID controller and its first and second network interface controllers. 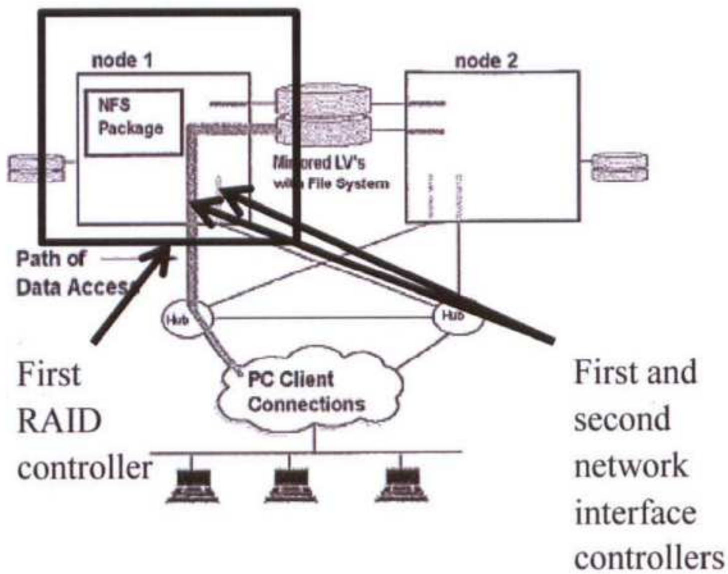 **Weygant, Fig. 4.1 (annotated)** The first network interface controller in the first RAID controller supplies data to the host computers by carrying file server requests. See, e.g., Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which one LAN card on each node is

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	active and the other is a standby. The active LAN carries file server requests from clients and also the cluster's own heartbeat messages."). Also, in the example of Fig. 4.1, the first network interface controller supplies data to the host computers by providing access to a Network File Serves (NFS) document system. Weygant at p. 112 ("This system uses highly available network file services (NFS). NFS is a general facility for accessing file systems remotely. In the example that follows, the NFS server software is made highly available, so that writers and editors do not lose access to their NFS mounted file systems for an extended period if the NFS server should fail. Figure 4.1 shows the basic configuration for this active/standby MC/ServiceGuard cluster."). Thus, Weygant discloses the first network interface controller in the first RAID controller supplies data to the host computers connected through the plurality of connection units.
[9.5] processes	Weygant and ServiceGuard render this feature obvious.

Claim Language of U.S. Patent No. 6,978,346	**Relevant Disclosure in Weygant, Mylex, and ServiceGuard**
information transmitted from the second network interface controller in the second RAID controller,	For instance, Weygant teaches that the LAN cards (network interface controllers) send heartbeat signals among the nodes (RAID controllers). Therefore, the network interface controllers in the RAID controllers process information transmitted from other network interface controllers in other RAID controllers. • See Weygant at p. 60 ("In a cluster, the high availability software establishes a communication link known as a heartbeat among all the nodes in the cluster on a subnet known as the heartbeat subnet. These messages allow the high availability software to tell if one or more nodes has failed."), disclosing that the nodes send heartbeat signals to each other. • See also Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which one LAN card on each node is active and the other is a standby. The active LAN carries file server requests from clients and also the cluster's own heartbeat messages."), disclosing that the heartbeats are transmitted from an active LAN interface (network

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	interface controller) on a node to another active LAN interface on another node. • See also Weygant at p. 123-124 ("An Ethernet configuration will be used, including two LAN interfaces per node attached to different hubs.... Data and heartbeats will use one LAN interface, and an RS232 connection between the two nodes will serve as a heartbeat backup in case of heavy user traffic on the LAN. The second LAN interface will serve as a standby."), disclosing that the heartbeat signals from one node to another are transmitted by the LAN interfaces (network interface controllers). Furthermore, Weygant goes into more detail by explaining which network interface controllers process information from which other network interface controllers. Specifically, an active LAN interface (network interface controller) communicates with an active LAN interface of another node (RAID controller). This concept is shown in Fig. 2.10 and is applicable to the examples of Figs. 4.1-4.4.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard

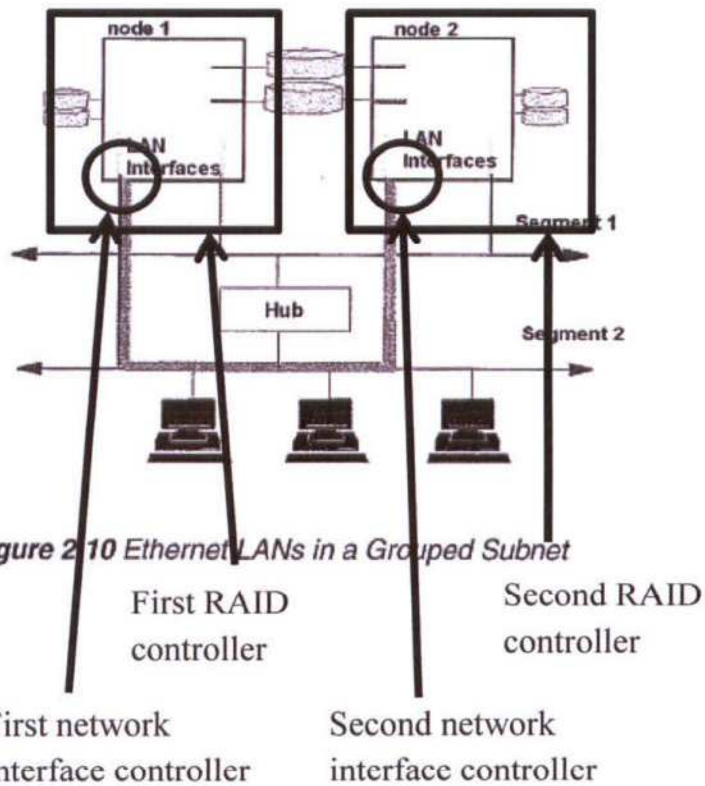

Weygant, Fig. 2.10 (annotated)

The determination of which LAN cards (network interface controllers) are active is simply a matter of configuration in the Weygant system. As explained above, ServiceGuard indicates that active and standby network interface controllers are determined at configuration and are, thus, merely a configuration

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	choice. See ServiceGuard at p. 70-71 ("It is recommended that you configure heartbeats on all subnets, including those to be used for client data. On the worksheet, enter the following for each LAN interface….IP Address: Enter this node's host IP address intended to be used on this interface. The IP address is a string of digits separated with periods in the form 'nnn.nnn.nnn.nnn'. If the interface is a standby and does not have an IP address, enter 'Standby.'). Therefore, when the first network interface controller in the first RAID controller is active and when the second network interface controller in the second RAID controller is active, the heartbeat signals are transmitted therebetween. Thus, Weygant and ServiceGuard disclose [the first network interface controller in the first RAID controller] processes information transmitted from the second network interface controller in the second RAID controller.
[9.6] wherein	Weygant and ServiceGuard render this feature obvious.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
the first network interface controller in the second RAID controller supplies data to the host computers connected through the plurality of connection units and	As noted above at [9.4], the LAN interfaces (network interface controllers) of the nodes (RAID controllers) supply data to the clients (plurality of host computers) by carrying file server requests and providing access to the NFS document system. In the analysis above at [9.4], node 1 (the first RAID controller) is running the NFS package and is supplying data to the host computers. Fig. 4.3 (reproduced below) shows that after remote switching, node 2 runs the NFS package and supplies data to the host computers. Thus, in Fig. 4.3, the first network interface controller in the second RAID controller supplies data to the host computers.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	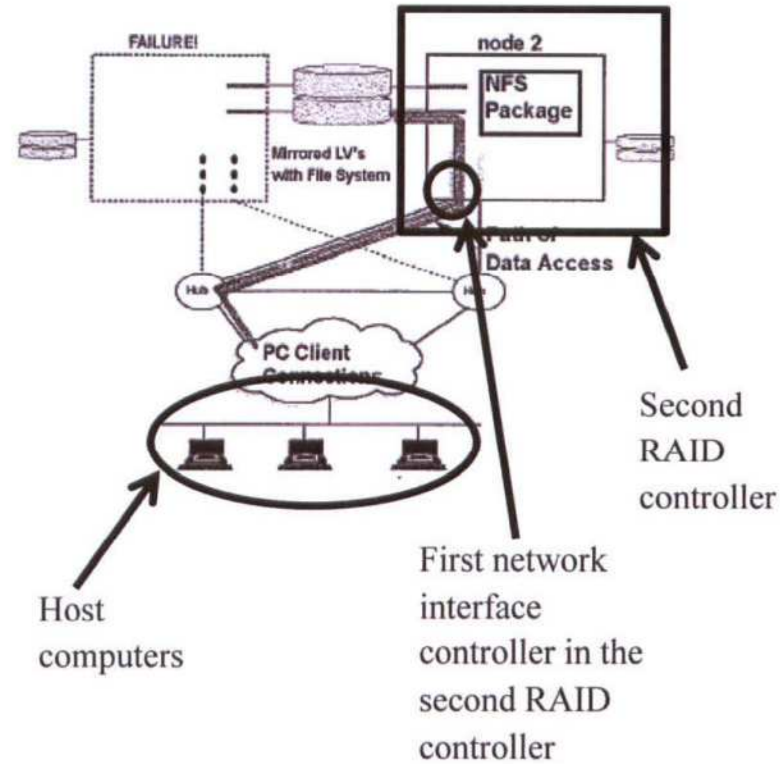 **Weygant, Fig. 4.3 (annotated)** The item that is characterized here as the "first network interface controller in the second RAID controller" is characterized as the "second network interface controller" at [9.5] above. However, designation as "first" or "second" is arbitrary because either LAN card can be active or standby, depending on configuration. But

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	for purposes of this example, the first network interface controller in the second RAID controller is the LAN interface shown as being active in Fig. 4.3. As shown above at [9.5], the determination of which LAN cards (network interface controllers) are active is simply a matter of configuration in the Weygant system, as evidenced by ServiceGuard. A person of ordinary skill in the art would have understood that a given LAN interface (network interface controller) at a node can be designated as active or standby. In other words, either one of the LAN interfaces (network interface controllers) in Figs. 4.1-4.4 can be the active LAN interface, depending on how the cluster is configured. Therefore, Weygant and ServiceGuard disclose the first network interface controller in the second RAID controller supplies data to the host computers connected through the plurality of connection units.
[9.7] processes	Weygant and ServiceGuard render this feature obvious.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
information transmitted from the second network interface controller in the first RAID controller,	As noted above at [9.5], the heartbeat signals from one node to another are transmitted by the LAN interfaces (network interface controllers). Also, as noted above at [9.5], either of the LAN interfaces (network interface controllers) at a node (RAID controller) may be configured as active, as evidenced by ServiceGuard. Therefore, when the first network interface controller in the second RAID controller is active and when the second network interface controller in the first RAID controller is active, the heartbeat signals are transmitted therebetween.
[9.8] wherein the second network interface controller in the first RAID controller is used for fault tolerance by	Weygant in view of ServiceGuard discloses this feature by disclosing fault-tolerant functionality, such that when a node fails, the other node runs the application package of the failed node. See, Weygant at p. 119 ("If the SPU experiences a failure, or if the operating system experiences a panic (fatal error), <u>the node will shut down, and MC/ServiceGuard on the other node will start the package in its alternate location</u>. The failover should take about 45 seconds in addition to the time required to start NFS.")

DHPN-1006

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
performing functions of the first network interface controller in the second RAID controller when the second RAID controller is faulty, and	See also Weygant at p. 125 ("In the event of SPU failure, the applications will continue running on the alternate node until the appropriate repair can be made on the failed node. After the loss of a node, of course, services will not be highly available until the repaired node re-enters the cluster."). Therefore, Weygant discloses that when a node (e.g., the second RAID controller) fails, the other node (the first RAID controller) runs the application package of the failed node. Cluster failover is illustrated in Weygant at Figs. 3.1 (before failover) and 3.2 (after failover), the principles of which are applicable to examples of Figs. 4.1-4.4.

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant, Mylex, and ServiceGuard</th></tr>
<tr><td></td><td>
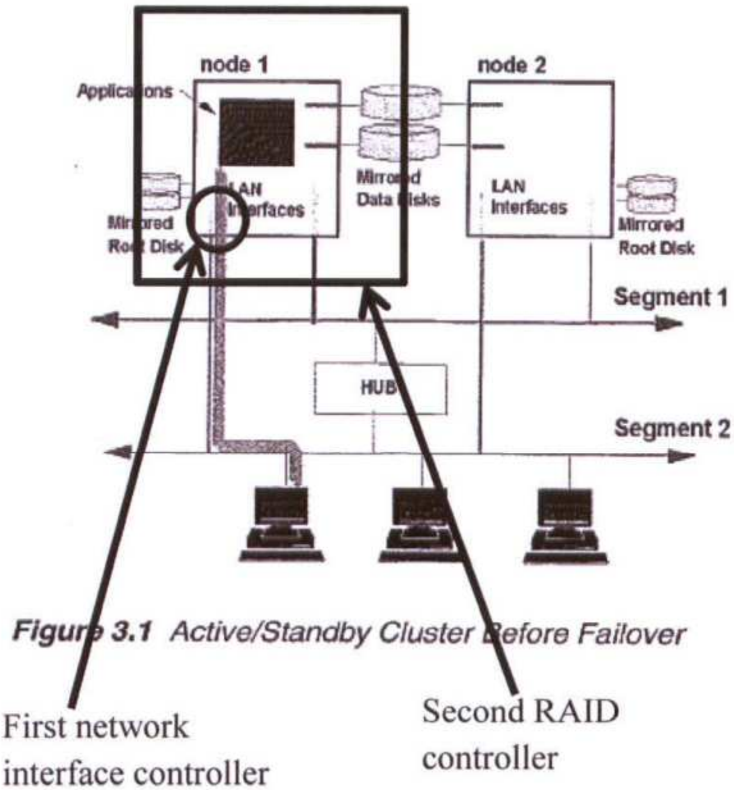

Figure 3.1 Active/Standby Cluster Before Failover

Weygant, Fig. 3.1 (annotated)
</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	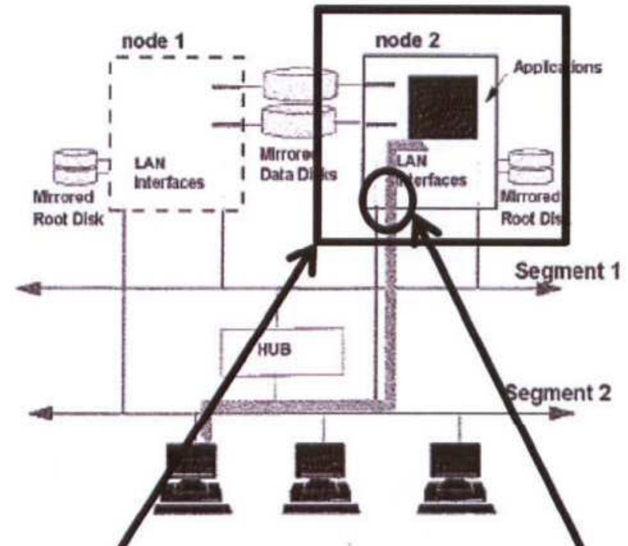 *Figure 3.2 Active/Standby Cluster After Failover* First RAID controller Second network interface controller **Weygant, Fig. 3.2 (annotated)** As shown above at Figs. 3.1 and 3.2, in a scenario in which the second network interface controller in the first RAID controller is active and in which the first network interface controller in the second RAID controller is active, a failure of the second RAID controller will result in the first RAID controller (and its active

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	network interface controller) taking over. Also, which network interface controllers are active is a matter of system configuration, as evidenced by ServiceGuard. I note that in these two figures, designations of "first RAID controller" and "second RAID controller" are reversed (left and right) with what is shown above when I refer to Fig. 4.1. That is because the failover concepts of Figs. 3.1 and 3.2 apply to the embodiments of Figs. 4.1-4.5, regardless of which node is designated as a "first" or a "second" RAID controller. In other words, the concept of failing over from one node to another applies just as well to one node as to the other. Thus, Weygant in view of ServiceGuard discloses the second network interface controller in the first RAID controller is used for fault tolerance by performing functions of the first network interface controller in the second RAID controller when the second RAID controller is faulty.
[9.9] wherein	Weygant in view of ServiceGuard discloses this feature by

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
the second network interface controller in the second RAID controller is used for fault tolerance by performing functions of the first network interface controller in the first RAID controller when the first	disclosing fault-tolerant functionality, such that when a node fails, the other node runs the application package of the failed node. As shown above at [9.8], Weygant discloses that when a node fails, the other node runs the application package of the failed node. See Weygant at p. 119 and 125. Cluster failover is illustrated in Weygant at Figs. 3.1 (before failover) and 3.2 (after failover), the principles of which are applicable to examples of Figs. 4.1-4.4.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
RAID controller is faulty, and	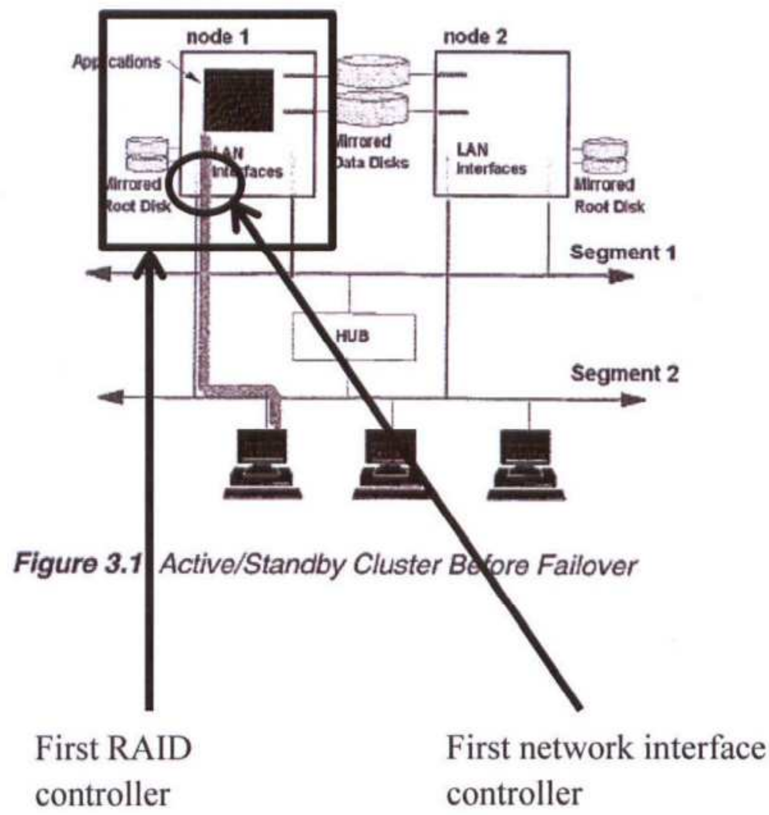 **Weygant, Fig. 3.1 (annotated)**

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard

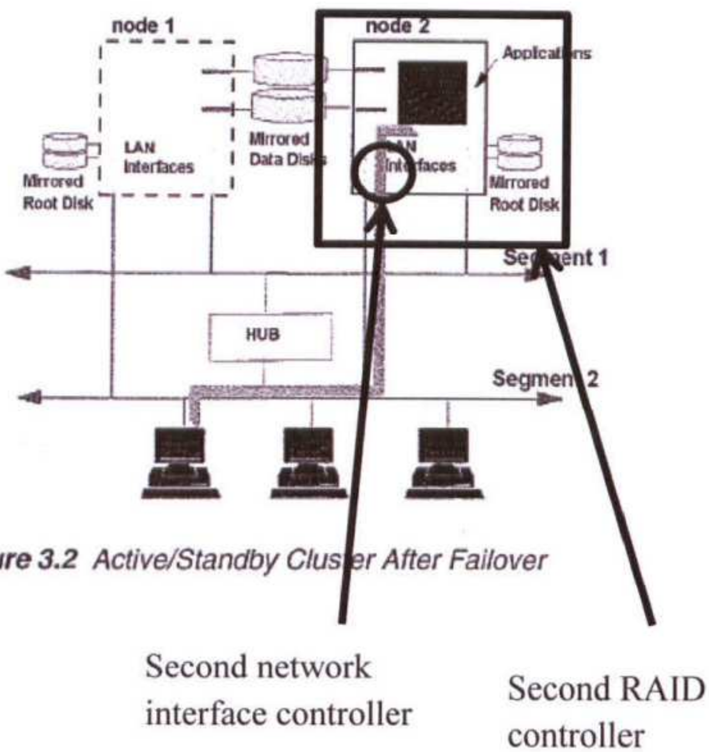

Figure 3.2 Active/Standby Cluster After Failover

Weygant, Fig. 3.2 (annotated)

As shown above at Figs. 3.1 and 3.2, in a scenario in which the second network interface controller in the second RAID controller is active and in which the first network interface controller in the first RAID controller is active, a failure of the first RAID controller will result in the second RAID controller (and its active

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	network interface controller) taking over. Also, which network interface controllers are active is a matter of system configuration, as evidenced by ServiceGuard. Thus, Weygant in view of ServiceGuard discloses the second network interface controller in the second RAID controller is used for fault tolerance by performing functions of the first network interface controller in the first RAID controller when the first RAID controller is faulty.
[9.10] wherein the first network controlling unit in the first RAID controlling unit exchanges information	Weygant discloses this feature by describing how the LAN interfaces (also referred to as LAN cards) exchange heartbeat signals. Weygant teaches that the nodes (RAID controllers) exchange heartbeat signals via their LAN interfaces (network interface controllers): • See Weygant at p. 60 ("In a cluster, the high availability software establishes a communication link known as a heartbeat among all the nodes in the cluster on a subnet

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
with the second network controlling unit in the second RAID controlling unit, and	known as the heartbeat subnet. These messages allow the high availability software to tell if one or more nodes has failed."), disclosing that the nodes send heartbeat signals to each other. • See also Weygant at p. 115 ("Figure 4.1 shows an Ethernet configuration in which <u>one LAN card on each node is active and the other is a standby</u>. The active LAN carries file server requests from clients <u>and also the cluster's own heartbeat messages</u>."), disclosing that the heartbeats are transmitted from an active LAN interface (network interface controller) on a node to another active LAN interface on another node. • See also Weygant at p. 123-124 ("An Ethernet configuration will be used, including two LAN interfaces per node attached to different hubs.... <u>Data and heartbeats will use one LAN interface</u>, and an RS232 connection between the two nodes will serve as a heartbeat backup in case of heavy user traffic on the LAN. The second LAN

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	interface will serve as a standby."), disclosing that the heartbeat signals from one node to another are transmitted by the LAN interfaces (network interface controllers). Furthermore, an active LAN interface (network interface controller) communicates with an active LAN interface of another node (RAID controller). This concept is shown in Fig. 2.10 and is applicable to the examples of Figs. 4.1-4.4.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	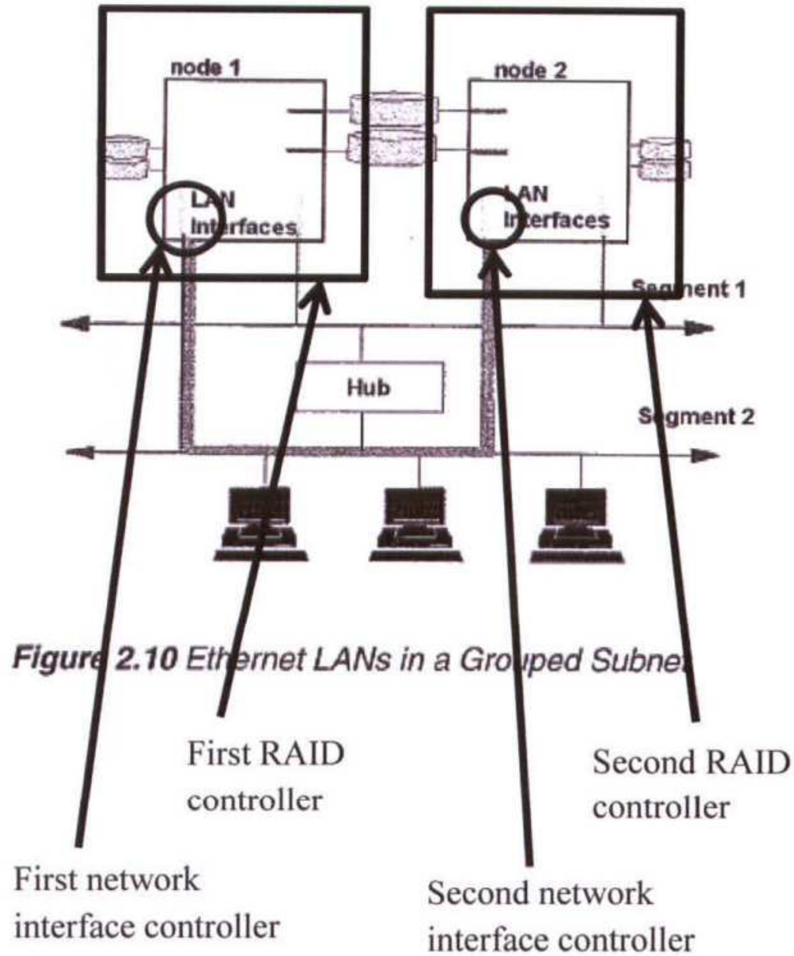 *Figure 2.10 Ethernet LANs in a Grouped Subnet* **Weygant, Fig. 2.10 (annotated)** In a scenario in which the first network interface controller in the first RAID controller is active and the second network interface controller in the second RAID controller is active, those two

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	network interface controllers would send and receive heartbeat signals. Thus, Weygant discloses the first network controlling unit in the first RAID controlling unit exchanges information with the second network controlling unit in the second RAID controlling unit.
[9.11] the second network controlling unit in the first RAID controlling unit exchanges information with the first network	Weygant discloses this feature by describing how the LAN interfaces (also referred to as LAN cards) exchange heartbeat signals. As described above at [9.10], Weygant teaches that the nodes (RAID controllers) exchange heartbeat signals via their LAN interfaces (network interface controllers). See Weygant at p. 60, 115, 123-124. Also, Weygant teaches an active LAN interface (network interface controller) communicates with an active LAN interface of another node (RAID controller). This concept is shown in Figs. 2.10 and 2.12 (reproduced below) and is applicable to the

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
controlling unit in the second RAID controlling unit.	examples of Figs. 4.1-4.4. 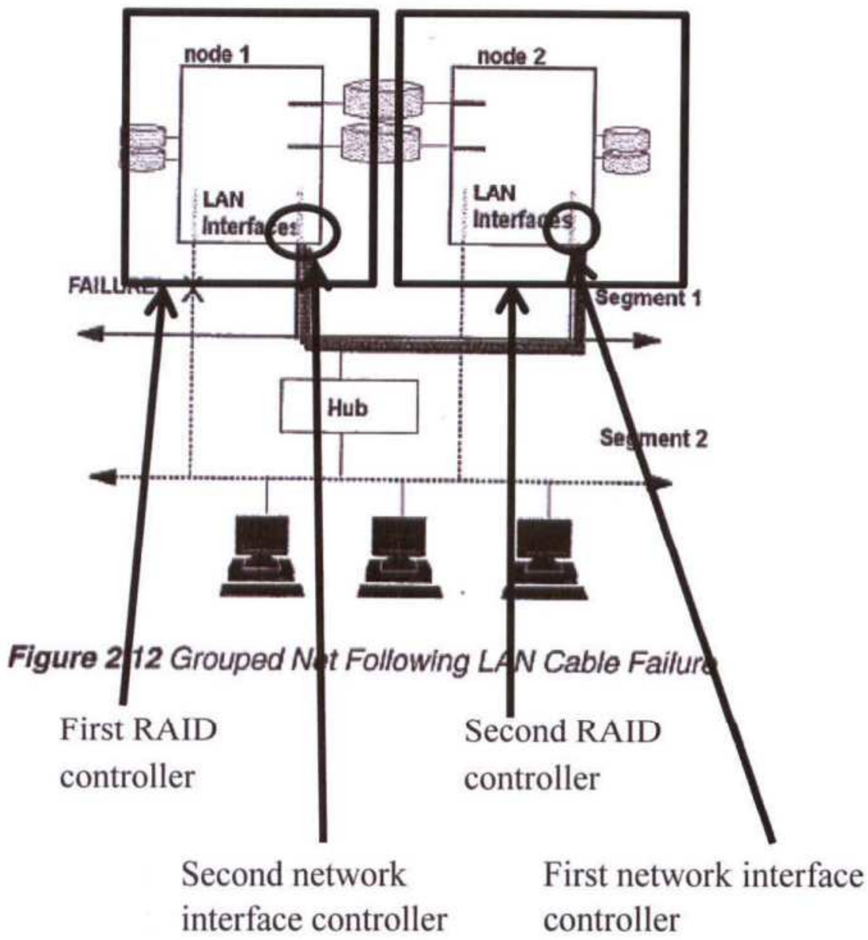 **Weygant, Fig. 2.12 (annotated)** In a scenario in which the second network interface controller in the first RAID controller is active and in which the first network interface controller in the second RAID controller is active, those

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ServiceGuard
	network interface controllers would send and receive heartbeat signals. Thus, Weygant discloses the second network controlling unit in the first RAID controlling unit exchanges information with the first network controlling unit in the second RAID controlling unit.

Challenge #4 - Claims 5-7 are obvious under 35 U.S.C. § 103(a) over Weygant and Mylex, further in view of ANSI

64. It is my opinion that Weygant and Mylex in further view of ANSI make claims 5-7 of the '346 patent obvious.

65. As shown above, Weygant and Mylex render obvious independent claim 1. Weygant discloses two hubs in its Fig. 4.1, where the hubs provide interconnections among the nodes and the PC clients. I do not believe that the term, "[X] of the at least [Y] connection ports is [are] coupled to one of the first network interface controlling unit and the third network controlling unit" requires that either [X] connection ports must be coupled to the first network controlling unit OR [X] connection ports must be coupled to the third network controlling unit. However, to the extent that someone might argue such a construction of the term above, I note that the workings of hubs and switches would have provided direct and indirect connections among all of the ports of a given hub or a switch so that a connection to a given hub or switch port serves as a connection to all of the devices coupled to the other ports on such hub or switch (thereby satisfying even such a construction). Hubs were old and well known in the art. In fact, the '346 patent admits that hubs including loop structures were old and known in the prior art. See '346 patent at 3:19-24. Also, ANSI is a technical document that describes the inner workings of a hub and specifically shows its loop structure in Fig. 1(a) (reproduced below in the chart). In the loop structure of ANSI, a connection to a given hub or switch port serves as a connection to all of the devices coupled to the other ports on such hub or switch.

66. It would have been obvious to a person of ordinary skill in the art in the 1999-2000 timeframe to implement the hubs of Weygant with a loop structure, such as the one shown in Fig. 1(a) of ANSI. A person of ordinary skill in the art would have made such a combination to ensure proper interconnection functionality at the hub. Additionally, applying the concepts of ANSI to the hubs of Weygant is merely a combination of prior art elements according to known methods to yield predictable results (the predictable result of providing network interconnections).

67. ANSI teaches use of a switch. ANSI at Fig. 1(b) (referencing a "fabric switch"). In fact, ANSI discusses the arbitrated loop as an alternative to a fabric switch. ANSI at p. 8 ("[t]his clause provides an overview of the structure, concepts, and mechanisms that allow two or more L_Ports to communicate without using a Fabric topology.").

68. ANSI also discloses that a fabric switch provides a topology so that a connection to a given port serves as a connection to all of the devices coupled to the other ports on such switch by virtue of fabric in the switch. ANSI at p. 9 ("Fabric topology may be configured to be non-blocking between any two

N_Ports. … A Fabric offers a way to take advantage of these natural pauses in communication, allowing fewer interconnects. The available bandwidth is shared between the N_Ports, but this sharing adds contention and therefore a management function is required . . . Because a Fabric topology may permit multiple paths between any two F_Ports in the Fabric (i.e., the meshing capability of the Fabric topology), a Fabric topology may be more robust.").

69. It would have been obvious to a person of ordinary skill in the art to have used network switch equipment or a switch to provide network interconnections in the system of Weygant as an added function to "take advantage of these natural pauses in communication" and because a switch fabric is "robust." ANSI at pp. 8-9. Also, using a network switch equipment or switch in place of Weygant's hubs is merely a simple substitution of one known element for another to obtain predictable results (the result of providing network interconnections).

70. The following table explains how Weygant, Mylex, and ANSI teach every element of claims 5-7 of the '346 patent.

Claim Language of U.S. Patent No. 6,978,346	**Relevant Disclosure in Weygant, Mylex, and ANSI**
[5.0] The apparatus as	See my analysis at claim 1 of the invalidity chart for Weygant and Mylex.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
recited in claim 1,	
[5.1] wherein said plurality of connecting units have at least three connection ports,	Weygant discloses this limitation. See, e.g., Weygant, Figure 4.1, showing each hub (a connecting unit) having at least four ports—one port in communication with node 1, one port in communication with node 2, one port to the other hub, and one port to the PC client connections. Since each hub has at least four ports, the total number of ports in the plurality of connecting units is at least eight.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
	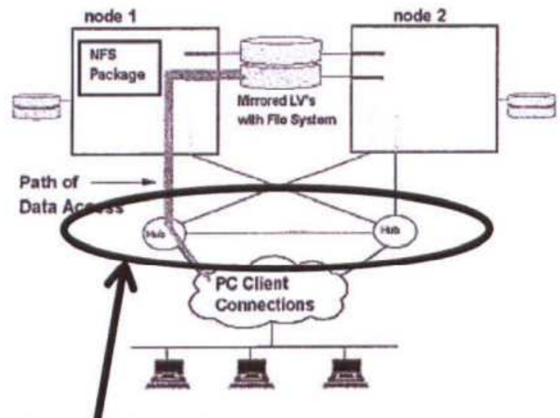 A plurality of connecting units, where each hub is shown with at least four ports **Weygant, Figure 4.1 (annotated)** Thus, Weygant's hubs with at least four ports each discloses said plurality of connecting units have at least three connection ports.
[5.2] two of the at least three connection ports is coupled to one	Weygant in view of ANSI renders obvious this limitation. Weygant teaches that its two hubs are coupled to each other, and that there are a total of eight connection ports. Weygant at Fig. 4.1. ANSI discloses that a hub connects each of its ports to all of the other ports on the hub. ANSI at p. 10 and Fig. 1(a). ANSI

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
of the first network interface controlling unit and the third network controlling unit	discloses the same single loop connecting each port as depicted in the Fig. 4 of the '346 patent, and such single loop may be implemented in the hubs of Weygant as well. 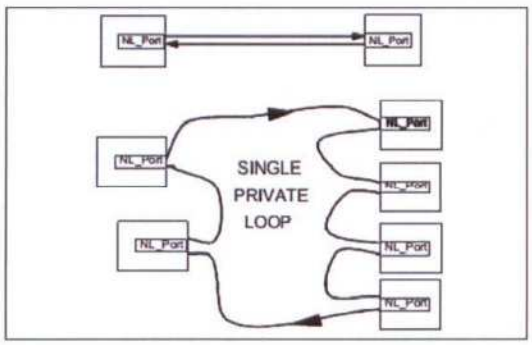 **ANSI, FIG. 1(a)** 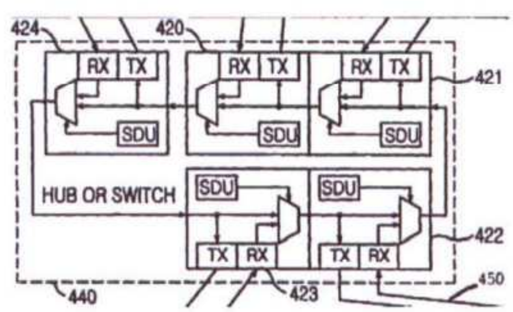 **'346 PATENT, FIG. 4** A person of ordinary skill in the art would have combined the concepts of Weygant and ANSI for the reasons I gave in the paragraphs above (e.g., to ensure proper interconnection

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
	functionality). Thus, as taught by Weygant and ANSI, in a given hub, a connection to a given port serves as a connection to all of the devices coupled to the other ports on such hub (by virtue of a loop structure in the hub). Thus, Weygant in view of ANSI discloses eight connection ports, in which each port is coupled, via a loop in the hub, to all of the devices on the network (including the respective LAN cards on the nodes). Thus, Weygant in view ANSI renders obvious two of the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit.
[5.3] and the rest of the connection ports being provided as a hub equipment connected	Weygant in view of ANSI renders obvious this limitation. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the clients (hosts). Also, ANSI discloses that a hub provides "a single private loop," so the rest of the ports are connected to the host computers, directly or indirectly. ANSI at p. 10 and Fig. 1(a). A person of ordinary skill in the art would have combined the teachings of Weygant and

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
with the numerous host computers.	ANSI for the reasons I gave in the paragraphs above (e.g., to ensure proper interconnection functionality). Thus, the ports provided by the hubs in Weygant in view of ANSI, where the hubs are connected with the PC clients, render obvious "the rest of the connection ports being provided as a hub equipment connected with the numerous host computers."
[6.0] The apparatus as recited in claim 1,	See my analysis at claim 1 of the invalidity chart for Weygant and Mylex.
[6.1] wherein said plurality of connecting units have at least three connection ports,	Weygant discloses this limitation. See, e.g., Weygant, Figure 4.1, showing each hub (a connecting unit) having at least four ports—one port in communication with node 1, one port in communication with node 2, one port to the other hub, and one port to the PC client connections. Since each hub has at least four ports, the total number of ports in the plurality of connecting units is at least eight.

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in Weygant, Mylex, and ANSI</th></tr>
<tr><td></td><td>
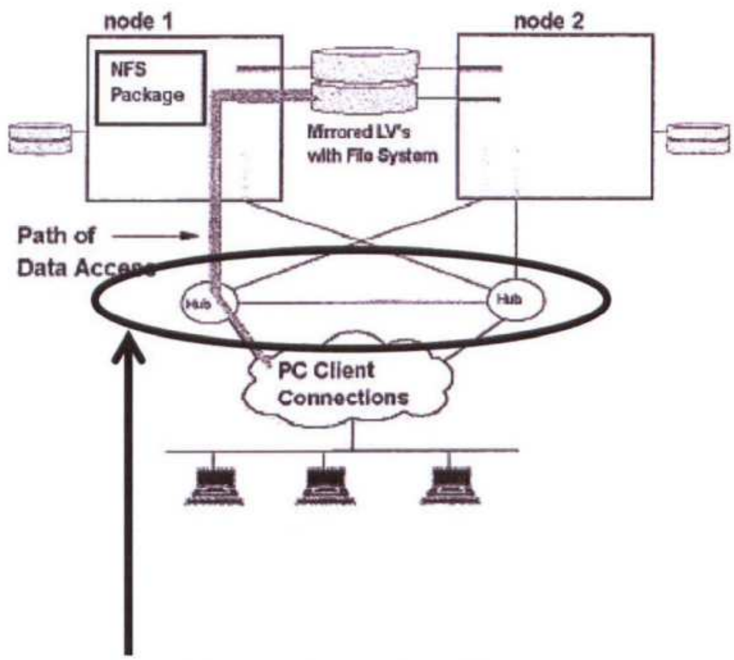

plurality of connection units where each hub is shown with at least four ports

Weygant, Figure 4.1 (annotated)

Thus, Weygant discloses said plurality of connecting units have at least three connection ports.
</td></tr>
<tr><td>[6.2] two of</td><td>Weygant in view of ANSI renders obvious this limitation.</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
the at least three connection port are coupled to one of the first network interface controlling unit and the third network controlling unit	Weygant teaches that its two hubs are coupled to each other, and that there are a total of eight connection ports. Weygant at Fig. 4.1. ANSI discloses that a switch is an alternative to a hub and that a switch connects each of its ports to all of the other ports on the switch. ANSI at pp. 8-9. ANSI discloses a switch that may be implemented in the system of Weygant. A person of ordinary skill in the art would have combined the concepts of Weygant and ANSI for the reasons I gave in the paragraphs above (e.g., to take advantage of natural pauses in communication and to be robust). Thus, as taught by Weygant and ANSI, in a given switch, a connection to a given port serves as a connection to all of the devices coupled to the other ports on such switch (by virtue of a fabric structure in the switch). Thus, Weygant in view of ANSI discloses eight connection ports, in which each port is coupled, via a fabric in the switch, to all of the devices on the network (including the respective LAN cards on

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
	the nodes). Thus, Weygant in view ANSI renders obvious two of the at least three connection port are coupled to one of the first network interface controlling unit and the third network controlling unit.
[6.3] and the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers.	Weygant in view of ANSI renders obvious this feature. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the clients (hosts). Also, ANSI discloses that a switch is an alternative to a hub and that a switch provides a topology so that the rest of the ports are connected to the host computers, directly or indirectly. ANSI at pp. 8-9. A person of ordinary skill in the art would have used network switch equipment in the system of Weygant at least for the reasons I gave in the paragraphs above as well (e.g., take advantage of natural pauses in communication and to be robust). Thus, Weygant in view of ANSI renders obvious the rest of the connection ports being provided as a network switch equipment

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
	connected with the numerous host computers.
[7.0] The apparatus as recited in claim 1,	See my analysis at claim 1 of the invalidity chart for Weygant and Mylex.
[7.1] wherein said plurality of connecting units have at least five connection ports,	Weygant discloses this limitation. See, e.g., Weygant, Figure 4.1, showing each hub (a connecting unit) having at least four ports—one port in communication with node 1, one port in communication with node 2, one port to the other hub, and one port to the PC client connections. Since each hub has at least four ports, the total number of ports in the plurality of connecting units is at least eight.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
	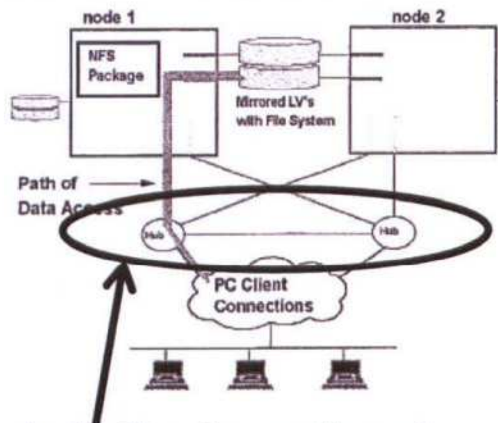 A plurality of connecting units, where each hub is shown with at least four ports **Weygant, Figure 4.1 (annotated)** Thus, Weygant's hubs with at least four ports each (eight ports total) discloses said plurality of connecting units have at least five connection ports.
[7.2] four of the at least five connection	Weygant in view of ANSI renders obvious this limitation. Weygant teaches that its two hubs are coupled to each other, and that there are a total of eight connection ports. Weygant at Fig. 4.1. ANSI discloses that a switch is an alternative to a hub and

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
ports is coupled to one of the first network interface controlling unit and the third network controlling unit	that a switch connects each of its ports to all of the other ports on the switch. ANSI at pp. 8-9. ANSI discloses a switch that may be implemented in the system of Weygant. A person of ordinary skill in the art would have combined the concepts of Weygant and ANSI for the reasons I gave in the paragraphs above (e.g., to take advantage of natural pauses in communication and to be robust). Thus, as taught by Weygant and ANSI, in a given switch, a connection to a given port serves as a connection to all of the devices coupled to the other ports on such switch (by virtue of a fabric in the hub). Thus, Weygant in view of ANSI discloses eight connection ports, in which each port is coupled, via a fabric in the switch, to all of the devices on the network (including the respective LAN cards on the nodes). Thus, Weygant in view ANSI renders obvious four of the at least five connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in Weygant, Mylex, and ANSI
[7.3] and the rest of the connection ports being provided as a switch connected with the numerous host computers.	Weygant in view of ANSI renders obvious this feature. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the clients (hosts). Also, ANSI discloses that a switch is an alternative to a hub and that a switch fabric provides a topology such that the rest of the ports are connected to the host computers, directly or indirectly. ANSI at p. 10 and Fig. 1(a). A person of ordinary skill in the art would have used a switch in the system of Weygant at least for the reasons I gave in the paragraphs above as well (e.g., to take advantage of natural pauses in communication and to be robust). Thus, Weygant in view of ANSI renders obvious the rest of the connection ports being provided as a switch connected with the numerous host computers.

Challenge #5 - Claims 1-3 and 5-8 are anticipated under 35 U.S.C. § 102(b) by the '950 patent

71. It is my opinion that the '950 patent anticipates claims 1-3 and 5-8 of the '346 patent.

72. The '950 patent is an IBM patent from the mid-1990s. The '950 patent discloses a "remote copy system" that provides data mirroring from one storage controller to another. See, e.g., '950 patent at Abstract and 7:37-40. For example, in the system of Fig. 3, a primary storage controller (e.g., storage controller 325) sends data for backup directly to a secondary storage controller (e.g., storage controller 335). '950 Patent at 7:37-40. The '950 patent proposes sophisticated techniques for the storage controllers to communicate with each other and with the host computers. The disclosure in the '950 patent of the storage controllers, the switches, and the communication paths among the various components teaches concepts that are claimed in claims 1-3 and 5-8 of the '346 patent.

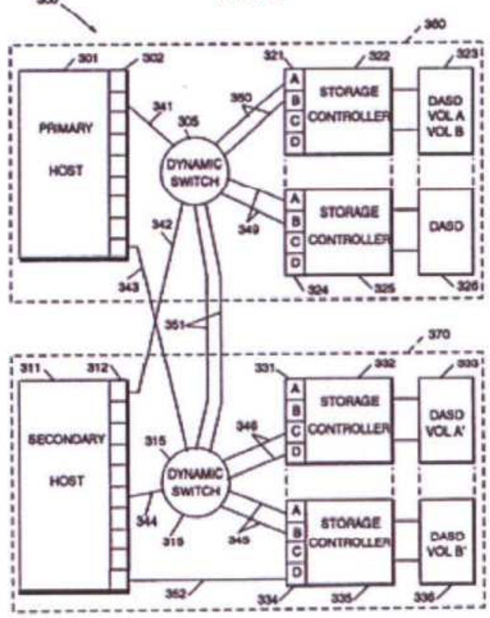

'950 Patent, Fig. 3

73. The following table explains how the '950 patent teaches every element of claims 1-3 and 5-8 of the '346 patent.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
[1.0] An apparatus for a redundant interconnection between multiple hosts and a RAID,	The '950 patent discloses this limitation. Fig. 3 of the '950 patent (reproduced and annotated below) illustrate primary and secondary hosts, which disclose the multiple hosts; the switches 305 and 315 as well as the connections thereto and therebetween disclose the redundant interconnections, and the DASDs disclose a RAID. The '950

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
comprising:	patent discloses that the DASDs can be implemented in a RAID configuration. '950 Patent at 2:5-11. FIG. 3 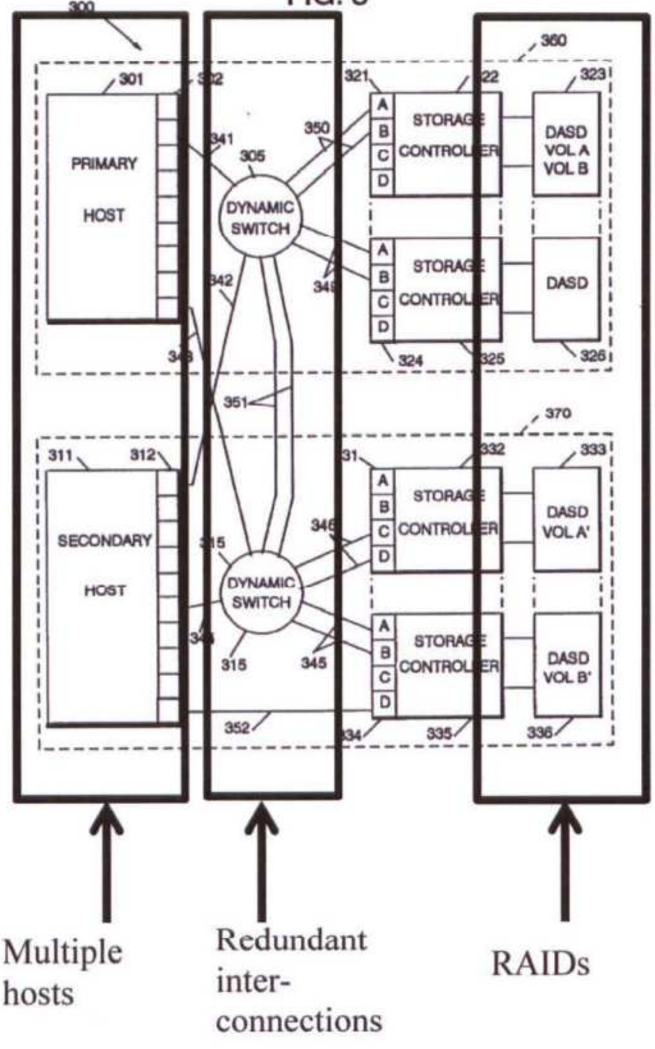 **'950 patent, Fig. 3 (annotated)**

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	Thus, the '950 patent discloses an apparatus for a redundant interconnection between multiple hosts and a RAID.
[1.1] a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers,	The '950 patent discloses storage controllers that teach the RAID controlling units. Fig. 3 of the '950 patent shows host computers 301 and 311 in communication with storage controllers (RAID controlling units) 322, 325, 332, 335. Also, the DASDs can be implemented in a RAID configuration. '950 Patent at 2:5-11.

<table>
<tr><th>Claim Language of U.S. Patent No. 6,978,346</th><th>Relevant Disclosure in the '950 Patent</th></tr>
<tr><td></td><td>

FIG. 3

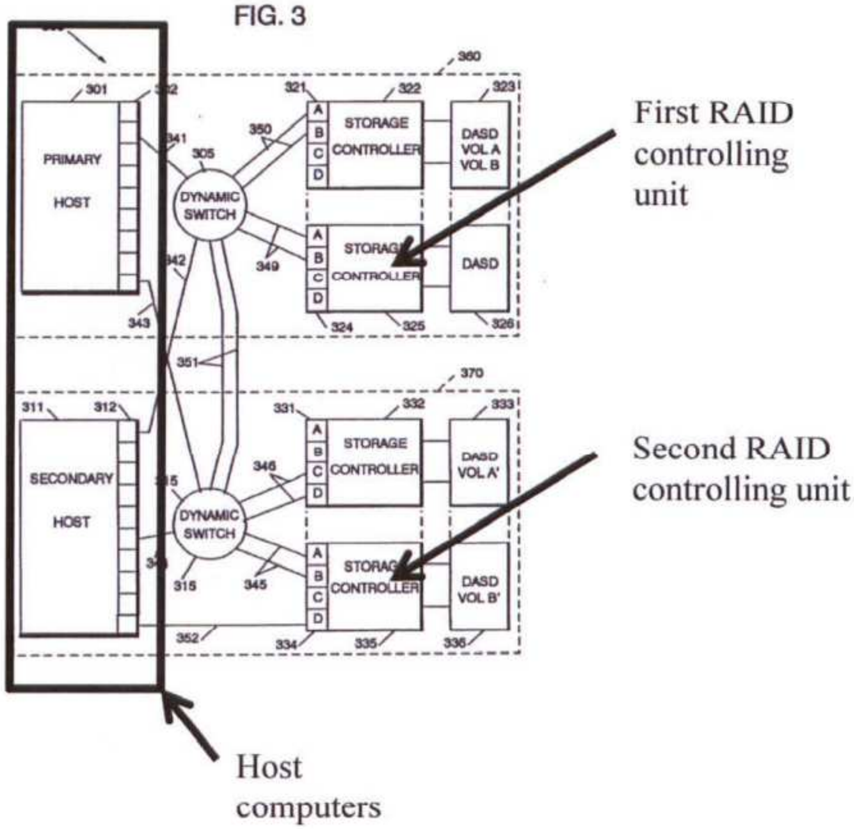

'950 Patent, Fig. 3 (annotated)

Furthermore, the '950 patent explains that the storage controllers 322, 325, 332, 335 process requirements of the host computers 301, 311 by providing storage. (The following Fig. 2 discussion applies equally well to the embodiment of Fig. 3):

</td></tr>
</table>

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	The primary host 201 can thus communicate with any secondary storage controller 232, 235, or the secondary host 211 via the dynamic switch 205 or 215. Likewise, the secondary host can communicate with any primary storage controller 222, 225, or the primary host 201 via the dynamic switch 205 or 215. Additionally, primary storage controllers 222, 225 can communicate with secondary storage controllers 232, 235, respectively. Thus, the primary host 201 could send data or records for back-up directly to the secondary storage subsystem (however, this may be undesirable due to the required primary host resources). More desirably, primary storage controllers 222, 225 send data or records for back-up directly to secondary storage controllers 232, 235, respectively. This communication is quicker since the primary host need only wait until the data or records are received in secondary storage controllers 232, 235 cache (see FIG. 1). '950 patent at 7:28-44. Thus, the '950 patent discloses a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
[1.2] the first RAID controlling unit including a first network controlling unit and a second network controlling unit,	The '950 patent discloses this limitation. The '950 patent at Fig. 3 has ports (ports A, B 324 teach first and second network controlling units, respectively; ports A, B 334 teach third and fourth network controlling units respectively). The four ports at each controller can be dynamically set to communicate either as a channel or control unit link-level facility. '950 Patent at 8:5-6.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	FIG. 3 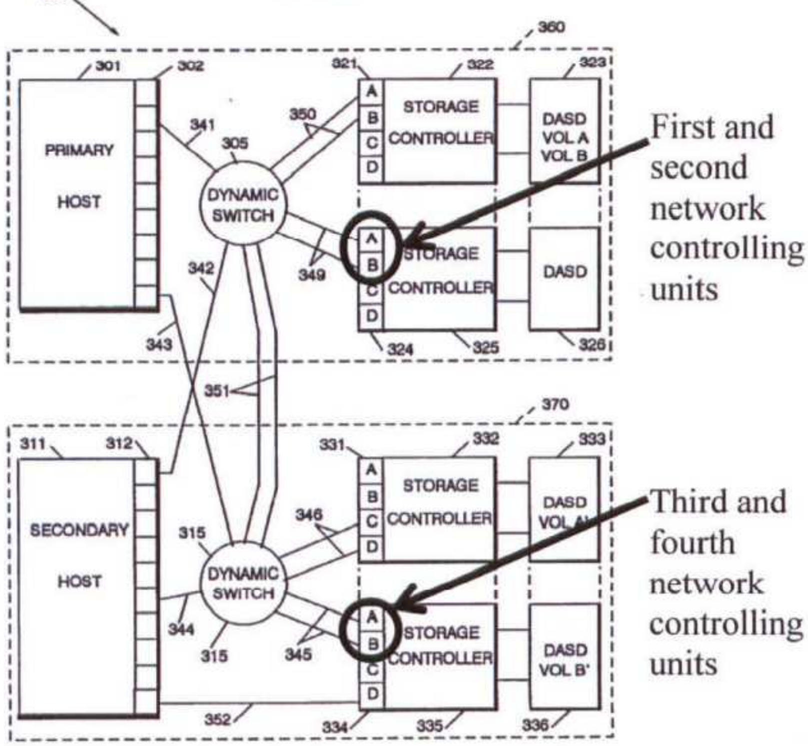 **'950 Patent, Figure 3 (annotated)** The examples of Figs. 4-8 discuss communications between storage controllers 325 and 335; thus, the following analysis addresses controllers 325 and 335 to be consistent with the examples in the '950 patent. However, it is noted that either of controllers 322 or 325 can be considered a first RAID controlling

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	unit, and either of controllers 332 or 335 can be considered a second RAID controlling unit for purposes of this analysis as well. Thus, the '950 patent's storage controllers including ports disclose the first RAID controlling unit including a first network controlling unit and a second network controlling unit.
[1.3] and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit;	See my analysis at [1.2] (above), showing third and fourth network controlling units. Thus, the '950 patent's storage controllers including ports disclose the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit.
[1.4] a	The '950 patent discloses this limitation.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers,	The '950 patent at Fig. 3, shows dynamic switches 305, 315. FIG. 3 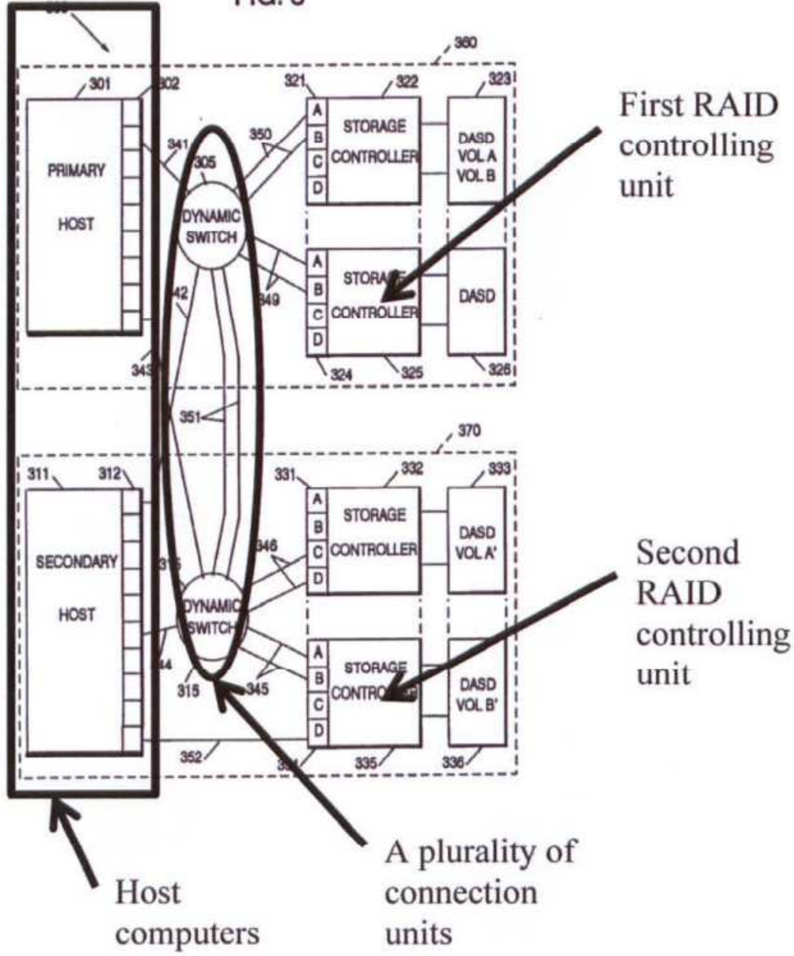 **'950 Patent, Fig. 3 (annotated)** Thus, the '950 patent's switches 305 and 315 disclose a plurality

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.
[1.5] wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting	The '950 patent discloses this limitation. For instance, the '950 patent discloses that the primary and secondary hosts can communicate with the primary and secondary storage controllers via the dynamic switches. This is evidenced by the following quote from the '950 patent, which applies equally well to both Figs. 2 and 3: The primary host 201 can thus communicate with any secondary storage controller 232, 235, or the secondary host 211 via the dynamic switch 205 or 215. Likewise, the secondary host can communicate with any primary storage controller 222, 225, or the primary host 201 via the dynamic switch 205 or 215. '950 patent at 7:28-35; see also 8:3-15 and Fig. 6, step 601. Thus, the '950 patent's communications between the hosts and the storage controllers discloses the first RAID controlling unit and

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
units,	the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units.
[1.6] and the first network controlling unit exchanges information with the fourth network controlling unit,	The '950 patent discloses this limitation. For instance, the '950 patent discloses that the ports of a primary storage controller (a first RAID controlling unit) communicate with the ports of a secondary storage controller (a second RAID controlling unit) via switches 305 and 315 (connection units). For example, the primary and secondary storage controller ports 321, 324, 331, and 334 can be dynamically set to communicate either as a channel or control unit link-level facility. Hence, primary storage controller 322, via port A 321, can communicate with primary host 301 by communication links 350, dynamic switch 305 and communication link 341, wherein port A 321 is a control unit link-level facility. Alternately, primary storage controller 322, via the same port A 321, can communicate with secondary storage controller 332 by communication links 350, dynamic switch 305, communication links 351, dynamic switch 315, and communication links 346, wherein

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	port A 321 acts as a channel link-level facility. '950 Patent at 8:3-15 (emphasis added). Examples of information exchange between the storage controllers 325, 335 include, *inter alia*, data mirroring (9:29-51), defining the peer-to-peer path (11:1-25), and establishing connection with second controller (11:60-12:15). The '950 patent also explains that when initiating a remote copy session (as in Fig. 4), ports A and B 334 (the third and fourth network controlling units) are used. Similarly, while performing the data mirroring of Fig. 5, ports A and B 324 (first and second network controlling units) are used. During this initialization process, ports 334 (A-B) have functioned in control unit link-level facility mode. ... At step 512, port 324 (A—B) is operating in channel link-level facility mode as will be described further.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	'950 Patent at 8:61-63 and 9:49-51. Thus, the '950 patent teaches that during normal operation, ports A and B 324 (first and second network control units) communicate with ports A and B 334 (third and fourth network control units). Furthermore, the '950 patent teaches that "the communication links between primary and secondary processors and between primary and secondary storage controllers may vary." '950 Patent at 13:13-16. Thus, the '950 patent teaches that the first network control unit (e.g., port A 324) exchanges information with either or both of the third and fourth network control units (ports A and/or B 334). Also, the second network control unit (e.g., port B 324) exchanges information with either or both of the third and fourth network control units (ports A and/or B 334). Accordingly, the '950 patent's communications between the storage controllers teaches the first network controlling unit exchanges information with the fourth network controlling unit.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
[1.7] and the second network controlling unit exchanges information with the third network controlling unit.	See my analysis at [1.6], showing that the second network control unit (e.g., port B 324) exchanges information with either or both of the third and fourth network control units (ports A and/or B 334). Accordingly, the '950 patent's communications between the storage controllers teaches the second network controlling unit exchanges information with the third network controlling unit.
[2.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[2.1] wherein said respective RAID	The '950 patent discloses this feature. See my analysis at [1.4] and Figure 3, showing the storage controllers 322, 325, 332, 335 (RAID controlling units) connected

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
controlling units are connected to the plurality of individual connecting units.	to the dynamic switches 305 and 315 (plurality of connecting units). Thus, the '950 patent's storage controllers being connected to switches discloses said respective RAID controlling units are connected to the plurality of individual connecting units.
[3.0] The apparatus as recited in claim 2,	See my analysis at claim 2.
[3.1] wherein the first network interface controlling unit is coupled to the	The '950 patent discloses this limitation. For instance, the '950 patent, Figure 3, shows ports A and B 324 (first and second network controlling units) both connected via links 349 to dynamic switch 305 (connecting unit) on an upper side and also both connected to dynamic switch 315 (connecting unit) on a lower side via links 349, switch 305, and links 351.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.	FIG. 3 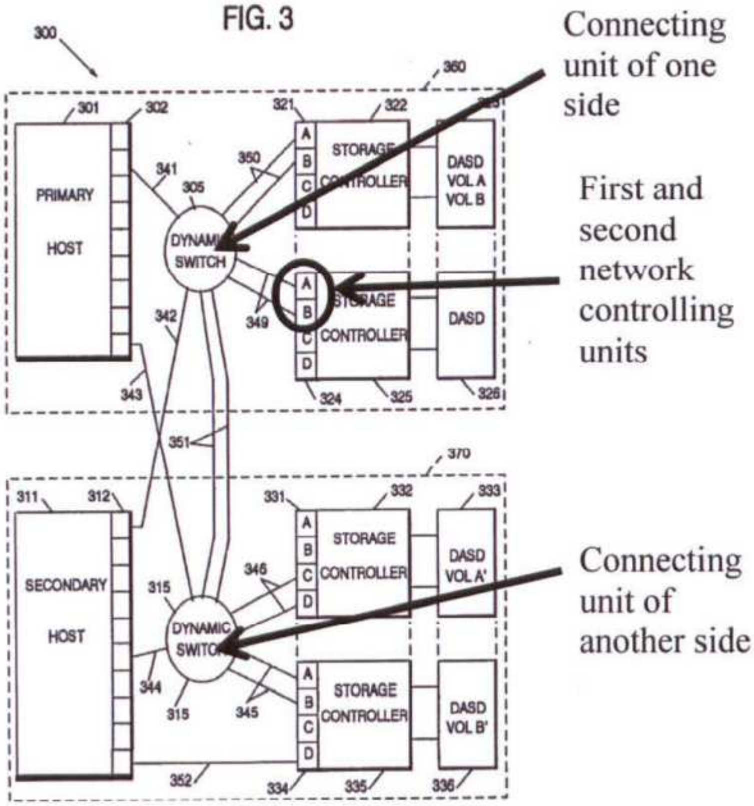 **'950 Patent, Figure 3 (annotated)** Thus, the '950 patents switches are each connecting units of different sides, and the ports A and B of storage controller 325 are network interface controlling units coupled thereto; thus, the '950 patent discloses the first network interface controlling unit is

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.
[5.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[5.1] wherein said plurality of connecting units have at least three connection ports,	The '950 patent discloses this limitation. The '950 patent, Figure 3, shows dynamic switch 305 having at least eight ports—one port for link 341, two ports for links 350, two ports for links 349, two ports for links 351, and one port for link 342. Dynamic switch 315 has a similar number of connection ports. The total number of connection ports is at least sixteen.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	FIG. 3 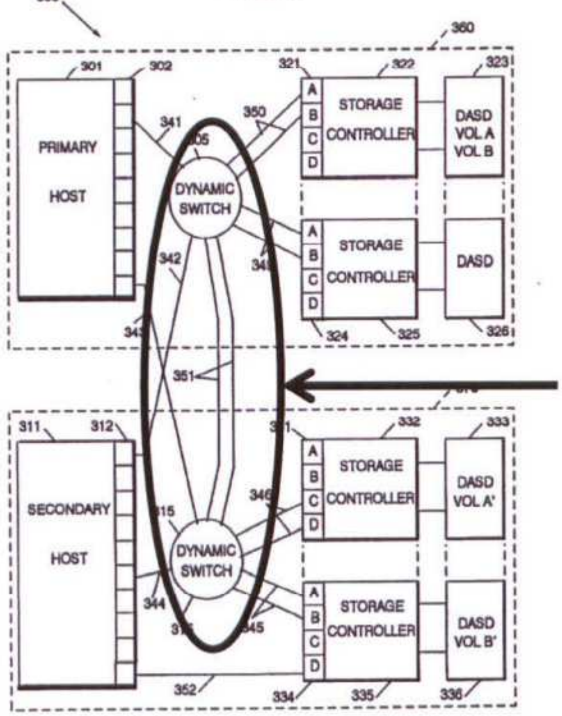 **'950 Patent, Figure 3 (annotated)**

DHPN-1006

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	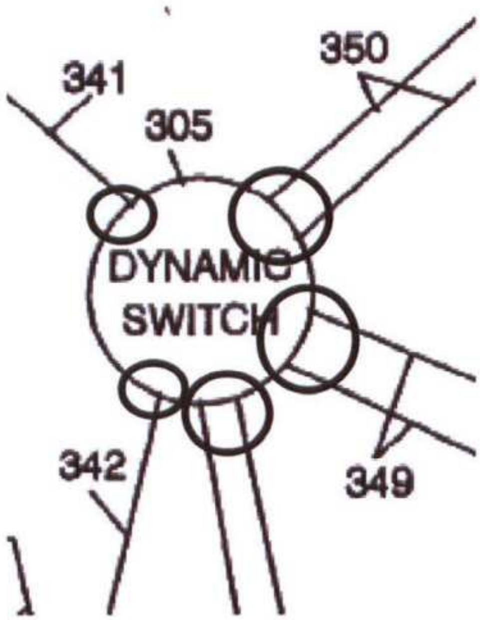 Switch 305 with eight ports circled. **'950 Patent, Figure 3 (truncated and annotated)** Thus, the '950 patent's ports at switch 305 disclose said plurality of connecting units have at least three connection ports; the '950 patent's ports at switch 315 also disclose this feature as well.
[5.2] two of	The '950 patent discloses this limitation.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit	The '950 patent, Figure 3 (truncated and annotated) at [5.1], shows one port of switch 305 connected to the first network interface controlling unit by link 349. The other seven ports of dynamic switch 305 are coupled to the first network interface controlling unit by virtue of being included in switch 305. (The '950 patent discloses that all ports of a switch are connected to all other ports of the switch by disclosing the various communication paths that include the switches. See '950 patent at 6:51-54, 7:10-13, 8:3-15.) Therefore, all of the at least eight ports of switch 305 are coupled to the first network controlling unit either directly or indirectly. In a similar way, all of the at least eight ports of dynamic switch 315 are coupled to the third network controlling unit either directly or indirectly. Thus, the '950 patent's ports at switch 305 discloses two of the at least three connection ports is coupled to … the first network interface controlling unit; similarly, the '950 patent's ports at switch 315 discloses two of the at least three connection ports is

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	coupled to … the third network controlling unit.
[5.3] and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers.	The '950 patent discloses this feature. Specifically, Figure 3 of the '950 patent shows dynamic switches 305 and 315. The '346 patent at 3:15-18 defines "hub" to include a hub or switch. Thus, dynamic switches 305 and 315 (plurality of connecting units) disclose the claimed hub equipment. The ports discussed above at [5.1] and [5.2] are thus provided as a hub equipment connected to the host computers 301 and 311. Thus, the '950 patent's switches 305 and 315 discloses the rest of the connection ports being provided as a hub equipment connected with the numerous host computers.
[6.0] The apparatus as recited in claim 1,	See my analysis at claim 1.
[6.1] wherein said plurality	The '950 patent discloses this limitation.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
of connecting units have at least three connection ports,	The '950 patent, Figure 3, shows dynamic switch 305 having at least eight ports—one port for link 341, two ports for links 350, two ports for links 349, two ports for links 351, and one port for link 342. Dynamic switch 315 has a similar number of connection ports. The total number of connection ports is at least sixteen. For further explanation, see my analysis at [5.1] (above), showing switch 305 enlarged with connection ports circled for emphasis.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	FIG. 3 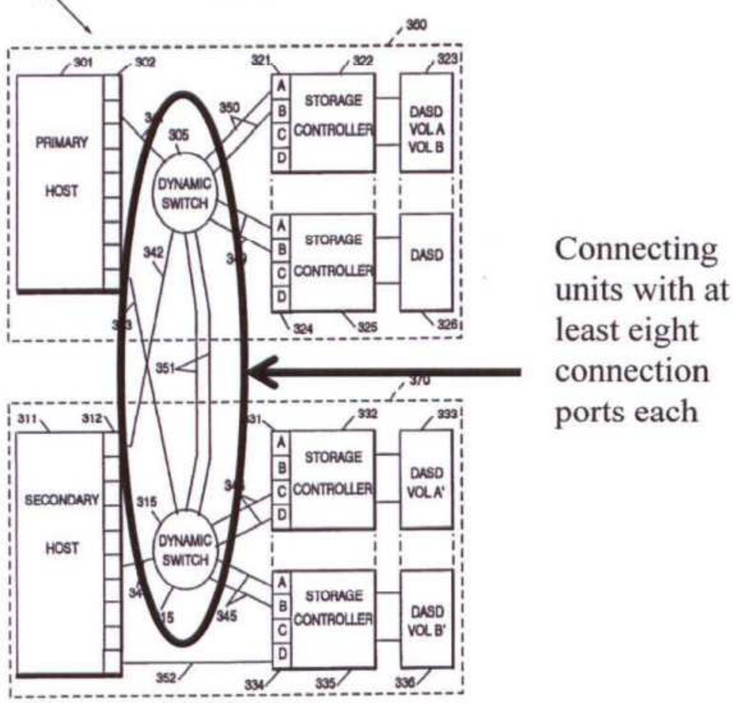 **'950 Patent, Figure 3 (annotated)** Thus, the '950 patent's ports at switches 305 and 315 disclose said plurality of connecting units have at least three connection ports.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
[6.2] two of the at least three connection port are coupled to one of the first network interface controlling unit and the third network controlling unit	The '950 patent discloses this limitation. The '950 patent, Figure 3 (truncated and annotated) at [5.1], shows one port of switch 305 connected to the first network interface controlling unit by link 349. The other seven ports of dynamic switch 305 are coupled to the first network interface controlling unit by virtue of being included in switch 305. (The '950 patent discloses that all ports of a switch are connected to all other ports of the switch by disclosing the various communication paths that include the switches. See '950 patent at 6:51-54, 7:10-13, 8:3-15.) Therefore, all of the at least eight ports of switch 305 are coupled to the first network controlling unit either directly or indirectly. In a similar way, all of the at least eight ports of dynamic switch 315 are coupled to the third network controlling unit either directly or indirectly. Thus, the '950 patent's ports at switch 305 discloses two of the at least three connection port are coupled to … the first network interface controlling unit; similarly, the '950 patent's ports at

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	switch 315 discloses two of the at least three connection port are coupled to … the third network controlling unit.
[6.3] and the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers.	The '950 patent discloses this limitation. Specifically, the '950 patent discloses dynamic switches 305 and 315 (plurality of connecting units) that are provided as network switch equipment. Thus, the '950 patent's switches disclose the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers.
[7.0] The apparatus as recited in	See my analysis at claim 1.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
claim 1,	
[7.1] wherein said plurality of connecting units have at least five connection ports,	The '950 patent discloses this limitation. The'950 patent, Figure 3, shows dynamic switch 305 having at least eight ports—one port for link 341, two ports for links 350, two ports for links 349, two ports for links 351, and one port for link 342. Dynamic switch 315 has a similar number of connection ports. The total number of connection ports is at least sixteen. For further explanation, see my analysis at [5.1] (above), showing switch 305 enlarged with connection ports circled for emphasis.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	FIG. 3 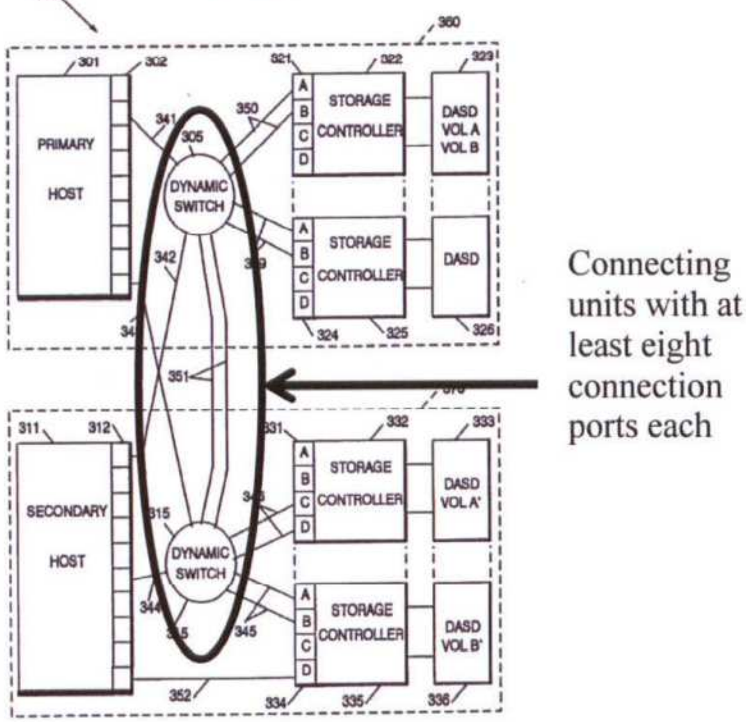 **'950 Patent, Figure 3 (annotated)** Thus, the '950 patent's switches discloses said plurality of connecting units have at least five connection ports.
[7.2] four of the at least	The '950 patent discloses this limitation. The '950 patent, Figure 3 (truncated and annotated) at [5.1],

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
five connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit	shows one port of switch 305 connected to the first network interface controlling unit by link 349. The other seven ports of dynamic switch 305 are coupled to the first network interface controlling unit by virtue of being included in switch 305. (The '950 patent discloses that all ports of a switch are connected to all other ports of the switch by disclosing the various communication paths that include the switches. See '950 patent at 6:51-54, 7:10-13, 8:3-15.) Therefore, all of the at least eight ports of switch 305 are coupled to the first network controlling unit either directly or indirectly. In a similar way, all of the at least eight ports of dynamic switch 315 are coupled to the third network controlling unit either directly or indirectly. Thus, the '950 patent's ports at switch 305 discloses four of the at least five connection ports is coupled to … the first network interface controlling unit; similarly, the '950 patent's ports at switch 315 discloses four of the at least five connection ports is coupled to … the third network controlling unit.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
	Thus, the '950 patent's switches disclose four of the at least five connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit.
[7.3] and the rest of the connection ports being provided as a switch connected with the numerous host computers.	The '950 patent discloses this limitation. Specifically, the '950 patent discloses dynamic switches 305 and 315 (plurality of connecting units) that are each provided as a switch. Thus, the '950 patent's switches disclose the rest of the connection ports being provided as a switch connected with the numerous host computers.
[8.0] The apparatus as recited in claim 1,	See my analysis at claim 1.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
[8.1] wherein the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second network interface controlling unit of said	The '950 patent discloses this limitation. For instance, as shown in Figure 3 (reproduced below): • Port A 324 (first network controlling unit) is connected to dynamic switch 305 (a first connecting unit) by link 349. • Port B 324 (second network controlling unit) is connected to the dynamic switch 315 (second connecting unit) by links 351, switch 305, and links 349. • Port A 334 (third network controlling unit) is connected to the dynamic switch 315 (second connecting unit) by link 345. • Port B 334 (fourth network controlling unit) is connected to switch 305 (first connecting unit) by link 345, switch 315, and links 351.

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
first RAID controlling unit being connected to a second connecting unit, the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit, and the	FIG. 3 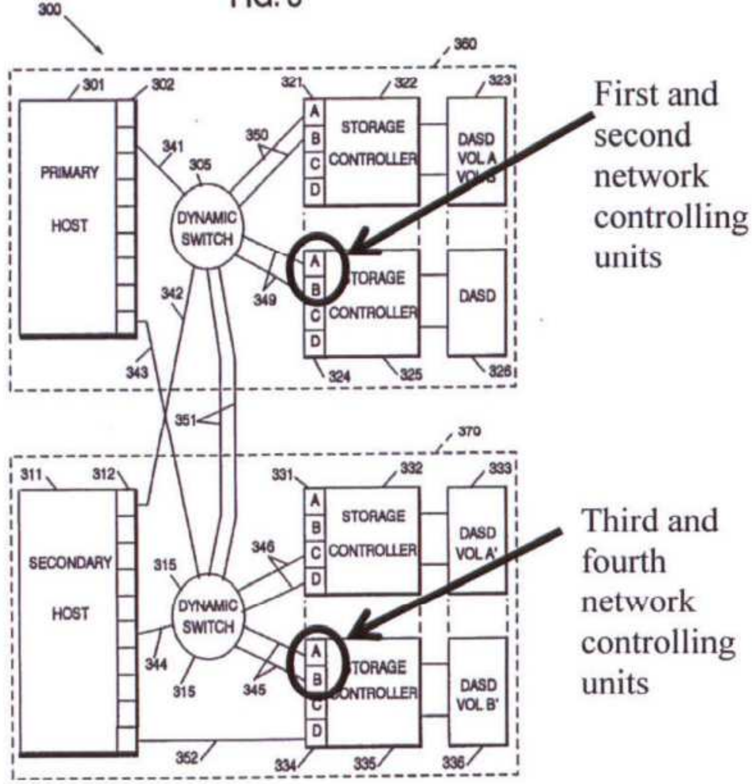 Thus, the '950 patent's connections between the ports of the storage controllers and the switches 305, 315 disclose the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit, the third network interface

Claim Language of U.S. Patent No. 6,978,346	Relevant Disclosure in the '950 Patent
fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.	controlling unit of the second RAID controlling unit being connected to the second connecting unit, and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.

Declaration

74. 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: September 26, 2013

By: Dr. M. Ray Mercer

Dr. M. Ray Mercer

MELVIN RAY MERCER

Independent Consultant
and
Professor Emeritus
Department of Electrical and Computer Engineering
Texas A & M University

BUSINESS ADDRESS:	311 Landsburg Lane College Station, TX 77845	**PHONE:** **CELL:** **FAX:** **EMAIL:**	(979) 485-0311 (979) 676-3109 (866) 769-1967 ray@rmercer.com
FEDEX ADDRESS: **SNAIL MAIL ADDRESS:**	Ray Mercer c/o The UPS Store Box # 333 3515-B Longmire Drive College Station, TX 77845	 **PHONE:**	 (979) 680-1410
UNIVERSITY ADDRESS:	Dept. of EE -- 214 Zachry Texas A & M University College Station, TX 77843 - 3259	**PHONE:** **FAX:** **EMAIL:**	(979) 862-3360 (979) 845-2630 mercer@ece.tamu.edu

EDUCATION:
The University of Texas at Austin, Ph.D., 1980
Stanford University, MS, 1971
Texas Tech University, BS, 1968

PROFESSIONAL REGISTRATION: Texas #62016

ACADEMIC POSITIONS:

Texas A & M University, Professor Emeritus, 2006
Texas A & M University, Professor and Chair Holder, 1/95-9/2005
University of Texas at Austin, Professor, 9/91-12/94
University of Texas at Austin, Associate Professor, 9/87-8/91
University of Texas at Austin, Assistant Professor, 1/83-8/87
University of Texas at San Antonio, Lecturer, 1977-80

OTHER PROFESSIONAL EXPERIENCE:

Independent Consultant for Mercer and Associates, College Station, TX 1984-present
Member of Technical Staff, Bell Laboratories, Murray Hill, NJ, 1980-83
Member of Technical Staff, Hewlett-Packard Laboratories, Palo Alto, CA, 1973-77
Research/Development Engineer, GTE Sylvania, Mountain View, CA, 1968-73

CONSULTING:

Hale and Dorr, Boston, Massachusetts, Intel, 2001
Fulbright and Jaworski, Houston, TX, Emachines, 1999-2000
Akin, Gump, Strauss, Hauer & Feld, 2000
Harris Corporation, Melbourne, FL, 1999-2001
SigmaTel, Austin, TX, 1999-2000
Martin Gruppen, Denmark, 1999
AT&T, Inc., Murray Hill, NJ, 1995-98

Sematech, Inc., Austin, TX 1994
Teradyne, Inc., Boston, MA, 1993-94
Integri-Test Corp., Commack, NY 1993
Spectrum Information Technologies, Dallas, TX 1993
Motorola Semiconductor, Austin, TX, 1987-88, 1991-99
Rockwell International, Newport Beach, CA, 1991, 1995
Teltech Resource Network, Minneapolis, MN, 1986-93
Cimflex Teknowledge, Pittsburgh, PA, 1989-90
IBM, Inc., Austin, TX, 1984, 1988-90
MCC, Austin, TX, 1989
CBS, New York, NY, 1985-86
Harris Data Communications Inc., Dallas, TX, 1984-86
Attorney General's Office, State of Texas, Austin, TX, 1984
Lockheed Missiles and Space Company, Austin, TX, 1983
Rothe Development Company, San Antonio, TX, 1979

TECHNOLOGY BASED BUSINESS EXPERIENCE:

Founder, Technical Advisor, and Owner of Conference Management Services, Inc., 1993 – 2013
Technical Advisory Board Member, Test Systems Strategies, Inc., Beaverton, OR, 1988-93
Advisor and Survey Developer, IBM, Inc., Austin, TX, 1988-90
Advisor, Early Cellular Telephone Minority Carriers, Lubbock, Texas, 1983 - 1987

HONORS AND AWARDS:

National and International:
Fellow of the Institute of Electrical and Electronics Engineers, 1994
National Science Foundation -- Presidential Young Investigator, 1986
Best Paper Award, VLSI Test Conference, Dana Point, CA, 1999
Best Paper Award, Design Automation Conference, San Francisco, CA, 1991
Best Paper Award, International Test Conference, Philadelphia, PA, 1982
Best Paper Award, Honorable Mention, Int. Test Conference, Washington, DC, 1988
Texas Tech Electrical Engineering Academy, Lubbock, Texas, 1999
Who's Who in America, 48th, 49th, 50th, 63rd, and 67th Editions, 1994, 1995, 1996, 2009, and 2012
Who's Who in America, Millennium Edition, 2000
Who's Who in America, 56th, 58th, and 59th Editions, 2002, 2004, and 2005
Who's Who in American Education, 3rd Edition, 1992-1993
Who's Who in America Finance and Industry, 31st Edition, 1999
Who's Who of Emerging Leaders in America, 7th Edition, 1990 & 1992
Who's Who in Science and Engineering, 3rd, 4th, 7th, 9th & 10th Editions, 2003-2009
Who's Who in the South and Southwest, 21st, 22nd, 23rd, 24th, 37th, 38th & 40th Editions, 1988, 1990, 1992, 1995, 2011, 2012, 2013 & 2014
Who's Who in the World, 10th, 11th, 17th, and 21st Editions, 1991, 1992, 2000 & 2004
Meritorious Service Award, IEEE Computer Society, 1993
Faculty Nominator and Advisor for Jennifer Dworak – Recipient of a National Science Foundation Graduate Research Fellowship 2000 – 2003
Co-Author and Advisor for Jennifer Dworak and Amy Wang – Recipient of the "IEEE Test Technology Technical Council Naveena Nagi Award for 2004" presented in Napa Valley, California

Local:
Computer Engineering Chair in Electrical Engineering, A&M, 1995-2005
Faculty Nominator and Advisor for Jennifer Dworak – Recipient of The Ethel Ashworth-Tsutsui Memorial Award for Graduate Student Research 2002
Texas A&M Outstanding Masters Thesis Award: Jennifer Dworak, 1999-2000
Listed in the Texas A&M Center for Teaching Excellence 2002 Eagle Award Booklet, May 3, 2002

Temple Foundation Endowed Professorship #3 in Engineering, UT, 1991-94
Engineering Foundation Endowed Faculty Fellowship in Engineering, UT, 1990-91
Werner W. Dornberger Centennial Teaching Fellowship in Engineering, UT, 1984-90
Engineering Foundation Faculty Award, UT, 1986
Outstanding Doctoral Dissertation: Honorable Mention, T. E. Kirkland, UT, 1986-87
MCC Sponsored Outstanding Student Paper Award: Bill Underwood, 1991-92
High School Valedictorian

PROFESSIONAL SOCIETIES AND ACTIVITIES:

Government:

National Science Foundation Advisory Committee for Microelectronic Information Processing Systems (MIPS), 1987-88
National Science Foundation Engineering Initiation Awards Evaluation Panel Member —Design, Tools and Test Program, 1987 and 1993
National Science Foundation Advisory Workshops
- Future of Testing and Design for Testability, June 30, 1989
- Future of VLSI and Computer-Aided Design, October 15-16, 1992

Presentation to the Texas State Board of Registration for Professional Engineers on Computer Engineering and suitable criteria for registration, 1985

Journals and Archival Publications:

Guest Editor, Special Issue on Design for Testability, *IEEE Design and Test of Computers*, October, 1986
Editor, Design for Testability, *IEEE Design and Test of Computers*, 1985-88
Guest Editor, *IEEE Transactions on Computer-Aided Design of Circuits and Systems*, 1988
Guest Editor, Special Issue on 1989 International Test Conference, *IEEE Design and Test of Computers*, April 1990.
Editorial Board Member, *Journal of Electronic Testing: Theory and Applications*, 1990-92
Editorial Advisory Board, *Microelectronics Journal: Circuits and Systems*, 2000 – 2003

Conferences and Workshops:

Finance Chair, Third IEEE Workshop on Microprocessor Test and Verification, (MTV'02), 2002
Program Committee, Ninth IEEE International Test Synthesis Workshop, (ITSW), 2002
Program Committee, Second IEEE Workshop on Microprocessor Test and Verification, (MTV 99), 1999
Exhibits Chairman, Fault-Tolerant Computing Symposium 1994
Planning Chairman, International Test Conference, 1992-93
Marketing Vice-Chairman, International Test Conference, 1990
Program Chairman, International Test Conference, 1989
Program Vice-Chairman, International Test Conference, 1988
Steering Committee, International Test Conference, 1987-93
Program Committee, International Test Conference, 1986-89
Program Committee, IEEE Design for Testability Workshop, 1988-96
Program Committee, International Conference on Computer-Aided Design, 1987
Program Committee, First MCC-University Research Symposium, Austin, TX, 1987

Local Offices:

Vice-Chairman, Central Texas Chapter, IEEE Computer Society, 1983-85
Chairman, Central Texas Chapter, IEEE Computer Society, 1985-86

Memberships:

Institute of Electrical and Electronics Engineers (IEEE), Fellow 1994, Life Member 2012

TEXAS A & M UNIVERSITY COMMITTEE ASSIGNMENTS:

City/County Committees:

Bryan/College Station Economic Development Group
- Marketing Committee for the Information Technology Task Force, 1999

University Committees:

Search Committee -- Associate Provost for Information Technology, 1997-99
Research Infrastructure Committee, 1998-99

College of Engineering Committees:

Computer Engineering Committee, Chairman, 2002
Tenure and Promotion Committee, 2000 - 2002
Computer Engineering Committee, Chairman, 1996-1999
Chair Holders Committee, 1995-
Compaq Liaison Committee 1996-
Computer Science Department Head Search Committee, 1996-98
ABET Review Committee, Computer Engineering, 1995
ABET Review Committee, Computer Engineering, Chair, 1998
Ad Hoc Committee to Study the Merger of the CS and EE Departments, 1996
PAM Advisory Committee, 1995-96
Spencer J. Buchanan Professorship Review Committee, 1997

Departmental Committees:

Computer Engineering Area Leader, 1995-present
Faculty Search Committee for the Computer Engineering Group, Chairman, 1995-present
Teaching Assignments for the Computer Engineering Group, 1995-present
Tenure and Promotion Committee, 1996-98, 2000-02, 2003-2005
Graduate Studies Committee, 1996-present
Faculty Advisory Committee, 1997-99
Strategic Planning Committee, 1998
Search Committee for the Eugene Webb Professorship, 1998-99
Search Committee for the Texas Instruments Jack Kilby Chair in Analog Engineering, 1998-99

UNIVERSITY OF TEXAS COMMITTEE ASSIGNMENTS:

University Committees:

Presentation to the MCC Site Selection Committee, the MCC Fact Finding Committee, and the MCC Technology Advisory Board, 1983
Science and Engineering Development Program Review for Dr. Thomas Everhart and Dr. Frank Press, National Academy of Sciences, 1984
Parking and Traffic Panel of the General Faculty, 1983-85
Hearing Officer for Faculty Grievances, 1987-88
University Council Representative, 1992-94
University Faculty Senate, 1992-94
Faculty Governance Committee, 1992-93

College of Engineering Committees:

Scholastic Appeals Committee, 1983-84
Ad Hoc Committee to Prepare a DOD Proposal for a Software Engr. Institute, 1984
State Agency Research Forum Speaker, May 10, 1984
Continuing Engineering Studies Committee, 1984-85
Ad Hoc Committee on Microelectronics and Computer Engineering, 1983-85
Presentation to Heads of State Agencies and Selected Federal Personnel, 1985
Computer Committee, 1985-86
GEC Faculty Meritorious Service Award Committee, 1987
Presentation to Industrial Representatives Research Forum, April 30, 1987
Undergraduate Degree Program Evaluation, 1986-88
Continuing Engineering Studies Committee, 1986-88
Televised Instruction Committee, 1987-91
Briefing for AT&T Visitors, November 21, 1991

Departmental Committees:

Committee on CAD/CAM and Advanced Graphics, 1983
Chairman, MCC Graduate Fellowships Recruiting Poster Committee, 1984
Microelectronics and Computer Engineering Research Support Committee, 1984
Chairman, Computation, Word Processing, and Telecom. Committee, 1984-85
Chairman, Industrial Liaison Committee, 1985-86

Computation, Word Processing, and Telecommunications Committee, 1985-86
Equipment Committee, 1984-86
ABET Accreditation for ECE in Computer Engineering (Site Visit), 1987
VLSI Course Area Committee, 1986-87
Chairman, Local Area Network Committee, 1987-88
Search Committee for New ECE Chairman, 1988-89
ECE Visiting Committee, 1989-90
Chairman, Annual Research Review Committee, 1988-92
Graduate Student Recruitment at Stanford University, January 1992
Alumni Committee 1992-93
Computer Engineering Research Center Executive Committee, 1988-93
Computer Sciences Liaison, 1988-93
Digital Systems Course Area Committee, 1988-93
Chairman, Teaching Effectiveness Committee, 1991-94
Budget Council, 1991-94
Computer Engineering Representative to the ECE Area Committee 1993-94
Junior Faculty Recruiting Committee, Computer Engineering, 1985-94

PUBLICATIONS:
Refereed Conference and Archival Journal Publications:

M. R. Mercer and V. D. Agrawal, "A Novel Clocking Technique for VLSI Circuit Testability," *IEEE Journal of Solid-State Circuits*, Vol. SC-19, April 1984, pp. 207-212.

K. S. Hwang and M. R. Mercer, "Derivation and Refinement of Fanout Constraints to Generate Tests in Combinational Logic Circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, October 1986, pp. 564-572.

T. Kirkland and M. R. Mercer, "Automatic Test Pattern Generation Algorithms," *IEEE Design and Test of Computers*, June 1988, pp. 43-55.

E. S. Park, M. R. Mercer, and T. W. Williams, "A Statistical Model for Delay-Fault Testing," *IEEE Design and Test of Computers*, February 1989, pp. 45-55. { NSF, ONR}

D. E. Ross, K. M. Butler, and M. R. Mercer, "Exact Ordered Binary Decision Diagram Size When Representing Classes of Symmetric Functions," *Journal of Electronic Testing: Theory and Applications*, vol. 2, no. 3, August 1991, pp. 243-259. { NSF}

E. S. Park, M. R. Mercer, and T. W. Williams, "The Total Delay Fault Model and Statistical Delay Fault Coverage," *IEEE Transactions on Computers*, vol. 41, no. 6, June 1992, pp. 688-698. { NSF, ONR}

E. S. Park and M. R. Mercer, "An Efficient Delay Test Generation System for Combinational Logic Circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 11, no. 7, July 1992, pp. 926-938. {NSF, ONR, TATP}

R. Kapur and M. R. Mercer, "Bounding Signal Probabilities for Testability Measurement Using Conditional Syndromes," *IEEE Transactions on Computers*, vol. 41, no. 12, December 1992, pp. 1580-1588. { NSF, ONR, SRC}

M. Heap and M. R. Mercer, "Least Upper Bounds on OBDD Sizes," *IEEE Transactions on Computers*, accepted for publication, July 1993.

C. Oh and M. R. Mercer, "Efficient Logic-Level Timing Analysis Using Constraint-Guided Critical Path Search," *IEEE Transactions on VLSI*, September 1996. {ONR}

J. Dworak, J. Wicker, S. Lee, M. R. Grimaila, K. M. Butler, B. Stewart, L-C. Wang, and M. R. Mercer, "Defect-Oriented Testing and Defective-Part-Level Prediction," IEEE *Design and Test of Computers*, January-February, 2001, Vol. 18, No. 1, pp. 31 - 41. {SRC, NSF, TATP}

M. R. Mercer, V. D. Agrawal and C. M. Roman, "Test Generation for Highly Sequential Scan-Testable Circuits through Logic Transformation," *International Test Conference 1981*, Philadelphia, PA, October 1981, pp. 561-565.

V. D. Agrawal and M. R. Mercer, "Testability Measures -- What Do They Tell Us?," *International Test Conference 1982*, Philadelphia, PA, November 1982, pp. 391-396. **(Best Paper of the 1982 ITC)**

M. R. Mercer and B. Underwood, "Correlating Testability with Fault Detection," *International Test Conference 1984*, Philadelphia, PA, October 1984, pp. 697-704.

E. Schell and M. R. Mercer, "CADTOOLS: A CAD Algorithm Development System," *The ACM/IEEE Design Automation Conference (22nd) Proceedings*, Las Vegas, NV, June 23-26, 1985, pp. 658-666.

J. Salick and M. R. Mercer, "Built-In Self Test Input Generator for Programmable Logic Arrays," *International Test Conference 1985*, Philadelphia, PA, November 1985, pp. 115-125.

K. S. Hwang and M. R. Mercer, "Derivation and Refinement of Fanout Constraints to Generate Tests in Combinational Logic Circuits," *IEEE International Conference on Computer-Aided Design*, Santa Clara, CA, November 1985, pp. 10-12.

K. S. Hwang and M. R. Mercer, "Informed Test Generation Guidance Using Partially Specified Fanout Constraints," *1986 International Test Conference*, Washington, DC, September 8, 1986, pp. 113-119.

R. K. Gaede, M. R. Mercer and B. Underwood, "Calculation of Greatest Lower Bounds Obtainable by the Cutting Algorithm," *1986 International Test Conference*, Washington, DC, September 9, 1986, pp. 498-505.

M. R. Mercer, "Logic Elements for Universally Testable Circuits," *1986 International Test Conference*, Washington, DC, September 9, 1986, pp. 493-497.

T. E. Kirkland and M. R. Mercer, "A Two Level Guidance Heuristic for ATPG," *IEEE Fall Joint Computer Conference*, Dallas, TX, November 2-6, 1986, pp. 841-847.

B. Underwood, J. Salick, M. R. Mercer and J. Kuban, "An Automatic Test Pattern Generation Algorithm for PLAs," *IEEE International Conference on Computer-Aided Design*, Santa Clara, CA, November 10-13, 1986, pp. 152-155.

T. Kirkland and M. R. Mercer, "A Topological Search Algorithm for ATPG," *The ACM/IEEE Design Automation Conference (24th) Proceedings,* Miami, FL, June 28-July 1, 1987, pp. 502-508.

S. P. Smith, M. R. Mercer and B. Brock, "Demand Driven Simulation: BACKSIM," *The ACM/IEEE Design Automation Conference (24th) Proceedings,* Miami, FL, June 28-July 1, 1987, pp. 181-187.

E. J. Aas and M. R. Mercer, "Algebraic and Structural Computation of Signal Probability and Fault Detectability in Combinational Circuits," *Proceedings of the 17th International Symposium on Fault-Tolerant Computing*, Pittsburgh, PA, July 6-8, 1987, pp. 72-77.

D. E. Ross and M. R. Mercer, "WAVE, A Concurrent Approach to Combinational Test Pattern Generation," *Proceedings of the MCC-University Research Symposium*, Austin, TX, July 14, 1987.

E. S. Park and M. R. Mercer, "Robust and Nonrobust Tests for Path Delay Faults of a Combinational Circuit," *Proc. 1987 International Test Conference,* Washington, DC, September 1-3, 1987, pp. 1027-1034.

S. P. Smith, B. Underwood and M. R. Mercer, "An Analysis of Several Approaches to Circuit Partitioning for Parallel Logic Simulation," *Proc. 1987 IEEE International Conference on Computer Design,* Rye Brook, NY, October 5-8, 1987, pp. 664-667.

C. T. Glover and M. R. Mercer, "A Method of Delay Fault Test Generation," *Proc. 25th ACM/IEEE Design Automation Conference,* Anaheim, CA, June 13-15, 1988, pp. 90-95.

R. K. Gaede, M. R. Mercer, K. M. Butler, and D. E. Ross, "CATAPULT: Concurrent Automatic Testing Allowing Parallelization and Using Limited Topology," *Proc. 25th ACM/IEEE Design Automation Conference,* Anaheim, CA, June 13-15, 1988, pp. 597-600.

E. S. Park, M. R. Mercer, and T. W. Williams, "Statistical Delay Fault Coverage and Defect Level for Delay Faults," *Proc. 1988 International Test Conference,* Washington, DC, September 12-14, 1988, pp. 492-499. **(Honorable Mention for Best Paper of the 1988 ITC)**

S. P. Smith, B. Underwood, and M. R. Mercer, "D3FS: Demand Driven Time First Deductive Fault Simulation," *Proc. 1988 International Test Conference,* Washington, DC, September 12-14, 1988, pp. 582-592.

C. T. Glover and M.R. Mercer, "A Deterministic Approach to Adjacency Testing for Delay Faults," *Proc. 26th ACM/IEEE Design Automation Conference,* Las Vegas, NV, June 25-29, 1989, pp. 351-356.

E. S. Park and M.R. Mercer, "An Efficient Delay Test Generation System for Combinational Logic Circuits," *Proc. 27th ACM/IEEE Design Automation Conference,* Orlando, FL, June 24-28, 1990, pp. 522-528.

K. M. Butler and M.R. Mercer, "The Influences of Fault Type and Topology on Fault Model Performance and the Implications to Test and Testable Design," *Proc. 27th ACM/IEEE Design Automation Conference,* Orlando, FL, June 24-28, 1990, pp. 673-678. {NSF, SRC}

D. E. Ross, K. M. Butler, R. Kapur, and M. R. Mercer, "Fast Functional Evaluation of Candidate OBDD Variable Orderings," *Proc. of The European Conference on Design Automation*, Amsterdam, The Netherlands, February 25-28, 1991, pp. 4-10. {NSF, ONR, SRC}

R. Kapur, K. M. Butler, D. E. Ross, and M. R. Mercer, "On Bridging Fault Controllability and Observability and Their Correlations to Detectability," Proc. *of The European Test Conference*, Munich, Germany, April 10-12, 1991, pp. 333-339. {NSF, ONR, SRC}

K. M. Butler and M. R. Mercer, "Quantifying Non-Target Defect Detection by Target Fault Test Sets," *Proc. of The European Test Conference*, Munich, Germany, April 10-12, 1991, pp. 91-100. {NSF, SRC}

T. W. Williams, B. Underwood, and M. R. Mercer, "The Interdependence Between Delay-Optimization of Synthesized Networks and Testing," *Proc. 28th ACM/IEEE Design Automation Conference*, San Francisco, California, June 17-19, 1991, pp. 87-92. **(Best Paper Award at 1991 DAC)** {none}

K. M. Butler, D. E. Ross, R. Kapur, and M. R. Mercer, "Heuristics to Compute Variable Orderings for Efficient Manipulation of Ordered Binary Decision Diagrams," *Proc. 28th ACM/IEEE Design Automation Conference*, San Francisco, California, June 17-19, 1991, pp. 417-420. {NSF, ONR, SRC}

E. S. Park, B. Underwood, T. W. Williams, and M. R. Mercer, "Delay Testing Quality in Timing-Optimized Designs," *Proc. 1991 International Test Conference,* Nashville, TN, October 28 - November 1, 1991, pp. 897-905. {NSF, ONR, TATP}

K. M. Butler, R. Kapur, D. E. Ross, and M. R. Mercer, "The Roles of Controllability and Observability in Design for Test," *Proc. 1992 IEEE VLSI Test Symposium*, Atlantic City, New Jersey, April 6-9, 1992. {NSF, ONR, SRC}

M. R. Mercer, R. Kapur, and D. E. Ross, "Functional Approaches to Generating Orderings for Efficient Symbolic Representations," *Proc. 29th ACM/IEEE Design Automation Conference*, Anaheim, California, June 9-11, 1992, pp. 624-627. {NSF, ONR, SRC}

R. Kapur, J. Park, and M. R. Mercer, "All Tests for a Fault are Not Equally Valuable for Defect Detection," *Proc. 1992 International Test Conference,* Baltimore, MD, September 20-24, 1992, pp. 762-769. {NSF, ONR, SRC}

M. A. Heap, W. A. Rogers, and M. R. Mercer, "A Synthesis Algorithm for Two-Level XOR Based Circuits," *Proc. IEEE International Conference on Computer Design,* Cambridge, MA, October 11-14, 1992, pp. 459-462. {TARP}

R. B. Brashear, D. R. Holberg, M. R. Mercer and L. Pillage, "ETA: Electrical-Level Timing Analysis," *IEEE International Conference on Computer-Aided Design*, Santa Clara, CA, November 8-12, 1992, pp. 258-262. {IBM, MOTO, ONR, SRC, TATP}

J. Park and M. R. Mercer, "An Efficient Symbolic Design Verification System," *Proc. IEEE International Conference on Computer Design,* Cambridge, MA, October 3-6, 1993, pp. 294-298. {SRC}

Eun Sei Park, and M. R. Mercer, "Switch-Level ATPG Using Constraint-Guided Line Justification," *Proc. 1993 International Test Conference,* Baltimore, MD, October 17-21, 1993, pp. 616-625. {none}

R. B. Brashear, N. Menezes, C. Oh, L. Pillage, and M. R. Mercer, "Predicting Circuit Performance Using Circuit-Level Statistical Timing Analysis," *Proc. of The European Design and Test Conference*, Paris, France, February 28-March 3, 1994. {ARPA, ONR, SRC}

J. Park, M. Naivar, R. Kapur, M. R. Mercer, and T. W. Williams, "Limitations in Predicting Defect Level Based on Stuck-at-Fault Coverage," *Proc. 1994 IEEE VLSI Test Symposium*, Cherry Hill, NJ, April 25-28, 1994, pp. 186-191. {ONR, SRC}

L-C Wang, M. R. Mercer, and T. W. Williams, "Enhanced Testing Performance via Unbiased Test Sets," *Proc. of The European Design and Test Conference*, Paris, France, March 6-9, 1995, pp. 294-302. {SRC}

J. Park, C. Oh, and M. R. Mercer, "Improved Sequential ATPG Using Functional Observation Information and New Justification Methods," *Proc. of The European Design and Test Conference*, Paris, France, March 6-9, 1995, pp. 262-266. {ARPA, SRC}

L-C. Wang, Sophia Kao, M. R. Mercer, and T. W. Williams, "On the Decline of Testing Efficiency as Fault Coverage Approaches 100%," *Proc. 1995 IEEE VLSI Test Symposium*, Princeton, NJ, April 30- May 3, 1995, pp. 74 - 83. {ONR, SRC}

C. Oh and M. R. Mercer, "Efficient Timing Analysis Using Constraint-Guided Critical Path Search," *Proc. Eighth Annual IEEE ASIC Conference and Exhibit,* Austin, TX, Sept. 18 - 20, 1995, pp. 289 - 293. {ARPA, ONR}

L-C. Wang, M. R. Mercer, and T. W. Williams, "On Efficiently and Reliably Achieving Low Defective Part Levels," *Proc. 1995 International Test Conference,* Washington, DC, October 23 - 25, 1995, pp. 616-625. {SRC, ONR}

T. W. Williams, R. Kapur, M. R. Mercer, R. H. Dennard, and W. Maly, "IDDQ Testing for High Performance CMOS -- The Next Ten Years," *Proc. of The European Design and Test Conference*, Paris, France, March 11-13, 1996, pp. 578-583. {none}

L-C. Wang and M. R. Mercer, "A Better ATPG Algorithm and Its Design Principles," *Proc. 1996 International Conference on Computer Design,* Austin, TX, October 7 - 9, 1996, pp. 248-253. {SRC}

J. Park and M. R. Mercer, "Using Functional Information and Strategy Switching in Sequential ATPG," *Proc. 1996 International Conference on Computer Design,* Austin, TX, October 7 - 9, 1996, pp. 254-260. {SRC}

L-C. Wang, M. R. Mercer, and T. W. Williams, "Using Target Faults to Detect Non-Target Defects," *Proc. 1996 International Test Conference,* Washington, DC, October 22 - 24, 1996, pp. 629-638. {SRC}

T. W. Williams, R. H. Dennard, R. Kapur, M. R. Mercer, and W. Maly, "IDDQ Test: Sensitivity Analysis of Scaling," *Proc. 1996 International Test Conference,* Washington, DC, October 22 - 24, 1996, pp. 786-792. {none}

M. R. Grimaila, S. Lee, J. Dworak, K. M. Butler, B. Stewart, H. Balachandran, B. Houchins, V. Mathur, J. Park, L-C. Wang, and M. R. Mercer, "REDO -- Random Excitation and Deterministic Observation -- First Commercial Experiment," *Proc. 1999 IEEE VLSI Test Symposium,* Dana Point, Calif., April 25 - 29, 1999, pp. 268-274. **(Best Paper Award at 1999 VLSI Test Symposium)** {TATP}

J. Dworak, M. R. Grimaila, S. Lee, L-C. Wang, and M. R. Mercer, "Modeling the Probability of Defect Excitation for a Commercial IC with Implications for Stuck-at Fault-Based ATPG Strategies," *Proc. 1999 International Test Conference,* Atlantic City, NJ, September 28 - 30, 1999, pp. 1031-1037. {TATP}

R. Mehler and M. R. Mercer, "Multi-level Logic Minimization Through Fault Dictionary Analysis," *Proceedings of the 1999 International Conference on Computer Design*, Austin, TX, October 10 - 13, 1999, pp. 315-318.

J. Dworak, M. R. Grimaila, S. Lee, L-C. Wang, and M. R. Mercer, "Enhanced DO-RE-ME Based Defect Level Prediction Using Defect Site Aggregation – MPG-D," *Proceedings of the 2000 International Test Conference,* Atlantic City, NJ, October 3 - 5, 2000, pp. 930-939. {TATP}

J. Dworak, M. R. Grimaila, B. Cobb, T-C. Wang, Li-C. Wang, and M. R. Mercer "On the Superiority of DO-RE-ME / MPG-D Over Stuck-at-Based Defective Part Level Prediction," *Proceedings of the Ninth Asian Test Symposium,* Taipei, Taiwan, December 4-6, 2000, pp. 151-157. {NSF, TATP}

T. W. Williams, M. R. Mercer, J. P. Mucha, and R. Kapur, "Code Coverage, What Does It Mean in Terms of Quality?" *Proceedings of the 2001 Annual Reliability and Maintainability Symposium,* Philadelphia, PA, January 22-25, 2001, pp. 420-424. {none}

S. Lee, B. Cobb, J. Dworak, M. R. Grimaila, and M. R. Mercer, "A New ATPG Algorithm to Limit Test Set Size and Achieve Multiple Detections of all Faults, *Proceedings of Design Automation and Test In Europe – DATE 2002,* Paris, France, March 4 – 8, 2002, pp. 94 - 99. {SRC, NSF}

J-J Liou, Li-C Wang, K-T Cheng, J. Dworak, M. R. Mercer, R. Kapur, and T. W. Williams, "Enhancing Test Efficiency for Delay Fault Testing Using Multiple-Clocked Schemes," *Proceedings of The 39th Design Automation Conference,* New Orleans, Louisiana, June 10 - 14, 2002, pp. 371 - 374.

J.-J. Liou, L.-C. Wang, K.-T. Cheng, J. Dworak, M. R. Mercer, R. Kapur, and T. W. Williams, "Analysis of Delay Test Effectiveness with a Multiple-Clock Scheme," *Proc. 2002 International Test Conference*, Baltimore, MD, October 8 - 10, 2002, pp. 407 - 416.

J. Dworak, J. Wingfield, B. Cobb, S. Lee, Li-C Wang, and M. R. Mercer, "Fortuitous Detection and its Impact on Test Set Sizes Using Stuck-at and Transition Faults," *Proceedings of the 2002 International Symposium on Defect and Fault Tolerance in VLSI Systems (DFT 2002),* Vancouver, Canada, November 6-8, 2002, pp. 177 - 185.

Li-C. Wang, A. Krstic, L. Lee, K-T. Cheng, M. R. Mercer, T. W. Williams, and M. S. Abadir, "Using Logic Models to Predict the Detection Behavior of Statistical Timing Defects," *Proceedings of the 2003 International Test Conference*, Charlotte, NC, September 30 - October 2, 2003, pp.1041 - 1050.

Y. Tian, M. R. Grimaila, W. Shi, and M. R. Mercer, "Minimizing Defect Levels Using a Linear Programming Based Optimal Test Selection Method," *Proceedings of the 2003 Asian Test Symposium*, Xi'an, P. R. China, November 17 - 19, 2003.

J. Wingfield, J. Dworak, and M. R. Mercer, "Function-Based Dynamic Compaction and its Impact on Test Set Sizes," *Proceedings of the 18th IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems,* Boston, MA, November 3 - 5, 2003, pp. 167 - 174.

J. Dworak, J. Wingfield, B. Cobb, and M. R. Mercer, "Balanced Excitation and its Effect on the Fortuitous Detection of Dynamic Defects," *Proceedings of Design Automation and Test In Europe – DATE 2004,* Paris, France, February 16 -20, 2004, pp. 1,066 – 1,071.

J. Dworak, D. Dorsey, A. Wang, and M. R. Mercer, "Excitation, Observation, and ELF-MD: Optimization Criteria for High Quality Test Sets," *Proceedings of the 2004 IEEE VLSI Test Symposium (VTS'04),* Napa Valley, CA, USA, April 25th - April 29th, 2004, pp. 9 - 15. **(IEEE Test Technology Technical Council Naveena Nagi Award for 2004)**

J. Dworak, J. Wingfield, and M. R. Mercer, "A Preliminary Investigation of Observation Diversity for Enhancing Fortuitous Detection of Defects," *Proceedings of the 19th IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems,* Cannes, France, October 11 - 13, 2004, pp. 460 - 468.

Chapters and Books:

K. M. Butler and M. R. Mercer, *Assessing Fault Model and Test Quality*, Kluwer Academic Publishers, 1991, ISBN 0 - 7923 - 9222 - 1.

V. D. Agrawal and M. R. Mercer, "Testability Measures -- What Do They Tell Us?," in *VLSI Testing and Validation Techniques*, IEEE Tutorial, H. Reghbati, editor, 1985, pp. 401-406.

Technical Reports:

M. R. Mercer and V. D. Agrawal, "Use of Clock Signal Redundancy for Testability," Bell Laboratories Technical Memorandum, July 1981.

C. M. Roman, V. D. Agrawal and M. R. Mercer, "An LSI Chip Designed for Testability," *Proceedings of the Bell System Conference on Electronic Testing*, Princeton, NJ, September 1981.

M. R. Mercer and V. D. Agrawal, "Applications for Testability Measures in VLSI Design," *Proceedings of the Bell System Conference on Electronic Testing*, Princeton, NJ, October 1982, pp. 52-58.

M. R. Mercer, "Computer Aided Design of Digital Systems," *Discovery -- Research and Scholarship at The University of Texas at Austin*, Vol. 9, No. 3, 1985, pp 17-21.

M. R. Mercer, "Testing and Design Verification of Electronic Components -- a Perspective of the Last 40 Years," *IEEE Computer, **(Invited Publication for the 40th Anniversary Issue)***, September, 1991.

Other publications:

J. Dworak, D. Dorsey, A. Wang, and M. R. Mercer with IBM Technical Contact M. W. Mehalic, "Estimating Mean Time to Failure in Digital Systems Using Manufacturing Defective Part Level," *4^{th} Annual IBM Austin Center for Advanced Studies Conference*, Austin, TX, February 21, 2003.

PROFESSIONAL SOCIETY PRESENTATIONS:

"Testability Strategies for Custom Polycell Designs," Computer Elements Workshop on VLSI Debug and Diagnosis, IEEE Computer Society, New York, NY, May 1982.

"Interpretations of Testability Measures," IEEE Design Automation Workshop, Michigan State University, East Lansing, MI, October 1982.

"Testability Measures -- What Do They Tell Us?," Automatic Testing and Measurement Exhibition, Wiesbaden, West Germany, March 1983 (by invitation as part of the "Best of Cherry Hill" Session).

"Testing Issues at the University of Texas," International Test Conference 1983, Philadelphia, PA, October 1983.

"Refinement of Statistical Evaluation of Testability Algorithms," (with B. Underwood), Seventh Annual IEEE Workshop on Design for Testability, Vail, CO, April 1984.

"SUBTLE -- A New Methodology for Structured Testability," Seventh Annual IEEE Workshop on Design for Testability, Vail, CO, April 1984.

"Why Calculating Observability is More Difficult than Controllability," Eighth Annual IEEE Workshop on Design for Testability, Vail, CO, April 1985.

"Automatic Test Pattern Generation for PLA's," (with J. Salick and B. Underwood), Fifth Annual IEEE West Coast Testing Workshop, Lake Tahoe, CA, April 1986.

"A Method for Empirical Evaluation of the Cutting Algorithm," (with R. Gaede), 9th Annual IEEE Workshop on Design for Testability, Vail, CO, May 1986.

"Exact Calculation of Fault Detection Probabilities in Multi-Output Combinational Circuits," (with E. Aas), Built-In Self-Test Workshop, Kiawah Island, Charleston, SC, March 11-13, 1987.

"Fault Model Comparisons and a Method for Testing with Vector Pairs," (with T. Glover), 10th Annual IEEE Workshop on Design for Testability, Vail, CO, April 23, 1987.

"A Review of Current Methods in Automatic Test Pattern Generation and Design for Testability," Nordic Workshop on Testing, Roros, Norway, March 15, 1988.

"An Empirical Comparison of Random-Pattern Testability under Two Classes of Delay Fault Coverage," (with T. Glover), 11th Annual IEEE Workshop on Design for Testability, Vail, CO, April 21, 1988.

"A Novel Segmentation Scheme for Pseudo-Exhaustive Testing," (with B. Stewart), 12th Annual IEEE Workshop on Design for Testability, Vail, CO, April 20, 1989.

"Distributed Demand-Driven Logic Simulation," (with S.P. Smith), International Workshop on CAD Accelerators, Oxford University, UK, September 21, 1989.

"Syndrome Estimation in Combinational Circuits Using Conditional Probabilities," (with R. Kapur), Built-In Self-Test Workshop, Kiawah Island, Charleston, SC, March 22, 1990.

"On Evaluating Target Fault Models and Non-Target Fault Detection," (with K. Butler), 13th Annual IEEE Workshop on Design for Testability, Vail, CO, April 17, 1990.

"Testing and Design Verification -- a Functional Perspective," **(invited plenary presentation)**, The International Conference on Computer Design, Cambridge, Mass., Sept. 17, 1990.

"Ordered Partial Decision Diagrams and their Applications," (with D. Ross), 14th Annual IEEE Workshop on Design for Testability, Vail, CO, April 17, 1991.

"Delay-Optimization of Synthesized Networks and its Impact on Testing," (with B. Underwood and T. W. Williams), 14th Annual IEEE Workshop on Design for Testability, Vail, CO, April 17, 1991.

"Enhanced Non-Target Defect Detection Based Upon Refined Test Sets for Target Faults," (with R. Kapur and J. Park), 15th Annual IEEE Workshop on Design for Testability, Vail, CO, April 23, 1992.

"A Comparison of Non-Target Defect Levels for Scanned and Non-Scanned Sequential Circuits When the Fault Coverage is 100%," (with J. Park and R. Kapur), 15th Annual IEEE Workshop on Design for Testability, Vail, CO, April 23, 1992.

"Testing and Design Verification -- a Functional Perspective," **(invited presentation)**, The Canadian Workshop on New Directions in Testing, Montreal, Quebec, Canada, May 21, 1992.

"Design for Testability and Built-In Self-Test -- Obstacles and Opportunities," **(invited Keynote)**, IEEE Workshop on Design for Testability and Built-In Self-Test, Vail, CO, April 20, 1994.

"Limitations in Predicting Defect Level Based on Stuck-at-Fault Coverage," (with J. Park, Mark Naiver, T. Williams, and R. Kapur), 15th Annual IEEE Workshop on Design for Testability, Vail, CO, April 20, 1994.

"Enhancing Testing Efficiency by Reducing Testing Biases," (with L-C. Wang, and T. W. Williams), IEEE Workshop on Design for Testability, Vail, CO, April 25, 1996.

"On Bridging Defects which Manifest as Delay Faults but are *NOT* IDDQ Testable," (with D. Ross, and G. Tu), IEEE Workshop on Design for Testability, Vail, CO, April 25, 1996.

"IDDQ Test: Sensitivity Analysis of Scaling," (with T. W. Williams, R. Kapur, R. Dennard, and W. Maly), IEEE Workshop on Design for Testability, Vail, CO, April 25, 1996.

"High Fault Coverage Behavioral Test Generation," (with L-C. Wang, and T. W. Williams), IEEE European Test Workshop, Montpelier, France, June 12 - 14, 1996.

"Failure Prediction Quality for Voltage versus IDDQ Testing Methods," (with R. Kapur and T. W. Williams), IEEE European Test Workshop, Cagliari (Grand Hotel Chia Laguna), Italy, May 28 - 30, 1997.

"Using Commercial ATPG Tools to Accurately Predict and Minimize Defective Part Level," (with J. Dworak, M. R. Grimaila, J. Wicker, K. M. Butler, B. Stewart, L-C. Wang, and T. W. Williams), Eighth International Test Synthesis Workshop, Santa Barbara, CA, March 26 - 28, 2001.

"A Study of Gate-Level Modeling Biases in DFT Methodologies for Testing Custom Designs," (with L-C. Wang, and M. S. Abadir), Eighth International Test Synthesis Workshop, Santa Barbara, CA, March 26 - 28, 2001.

"A Statistical Analysis of the Sensitivity to Defective Part Level Model Parameters during Test Pattern Set Selection (with J. Dworak, M. Grimaila, K. Butler, Jason Wicker and B. Stewart), The Ninth International Test Synthesis Workshop, Santa Barbara, CA, March 25 - 27, 2002. **(Best Student Presentation Award of the Ninth ITSW – student presenting was Jennifer Dworak)**

"The Effect of Uncertainty in the Model Parameter Tau on the Effectiveness of Test Sets Optimized with MPG-D," (with J. Dworak, M.R. Grimaila, J. Wingfield, B. Cobb, S. Lee, J. Wicker, K. Butler, B. Stewart, and B. Underwood), 3rd IEEE International Workshop on Microprocessor Test and Verification, Austin, TX, June 6-7, 2002.

"A New Estimator for Mean Time to First Failure: How Bad Were Those Defective IC's We Missed?" (with J. Dworak, D. Dorsey, and A. Wang," Tenth International Test Synthesis Workshop, Santa Barbara, CA, March 31-April 2, 2003.

"Evaluating a Greedy ATPG Algorithm for Generating Compact Transition Test Sets in Accordance with the Principles of DO-RE-ME," (with S. Lee, J. Dworak, and B. Cobb), 4th International Workshop on Microprocessor Test and Verification, Austin, TX, May 29-30, 2003.

"Binary Decision Diagrams and their Applications in Manufacture Testing," (with J. Wingfield, and B. Cobb), Eleventh International Test Synthesis Workshop, Santa Barbara, CA, April 5 - 7, 2004.

"Defect Delectability Classes and Their Effect on Optimal Test Pattern Generation Strategies," (with J. Dworak), Eleventh International Test Synthesis Workshop, Santa Barbara, CA, April 5 - 7, 2004.

"Reducing Structural Bias: An Initial Look at Observation Diversity," (with J. Dworak and J. Wingfield), Fifth International Workshop on Microprocessor Test and Verification, Austin, TX, September 8 - 10, 2004.

INVITED LECTURES:

"Computer-Aided Testing and Simulation," First Annual Research Review, Department of Electrical and Computer Engineering, The University of Texas at Austin, May 8, 1984.

"Automatic Test Pattern Generation for Digital Logic Circuits," Second Annual Research Review, Department of Electrical and Computer Engineering, The University of Texas at Austin, May 7, 1985.

"Computer-Aided Testing and Simulation," Texas Instruments, Dallas, TX, June 1985.

"New Directions in Logic Design for Testability," International Business Machines Corporation, Purchase, NY, April 8, 1986.

"New Directions in Logic Design for Testability," Semiconductor Research Corporation, Research Triangle Park, NC, April 11, 1986.

"Research in Logic Testing at the University of Texas at Austin," Weekly Undergraduate Seminar, Mississippi State University, Columbus, MS, November 6, 1986.

"New Issues in Design for Testability," Stanford University, Stanford, CA, November 18, 1986.

"New Issues in Design for Testability," Tektronix Research Laboratories, Beaverton, OR, November 19, 1986.

"Some New Results Using Structured Logic Design Methods," McGill University, Montreal, Quebec, Canada, March 10, 1987.

"A New Design for Testability Method," General Electric Central Research and Development Laboratories, Schenectady, NY, April 9, 1987.

"The Value of Endowed Funds for Research at The University of Texas at Austin," Endowed Donors Dinner, February 26, 1988.

"The Boolean Difference from a New Perspective," The Technical University of Trondheim, Trondheim, Norway, March 17, 1988.

"Automatic Test Pattern Generation for Digital Logic Circuits," IEEE Computer Society, The University of Texas at Austin, April 6, 1988.

"Automatic Test Pattern Generation for Digital Logic Circuits," Schlumberger Austin Systems Center, Austin, Texas, April 7, 1988.

"Automatic Test Pattern Generation for Digital Systems," AT&T, Murray Hill, New Jersey, April 15, 1988.

"An Empirical Comparison of Random-Pattern Testability Under Two Classes of Delay Fault Coverage," NCR Technical Information Exchange Session, MCC, Austin, Texas, May 4, 1988.

"Designing and Testing Integrated Circuits," The Honors Colloquium, The University of Texas at Austin, July 26, 1985, July 26, 1986, July 25, 1987, July 22, 1988, and July 22, 1989.

"Statistical Delay Fault Coverage and Defect Level for Delay Faults," IBM, Austin, Texas, September 22, 1988.

"Statistical Delay Fault Coverage and Defect Level for Delay Faults," Northeastern Univer sity, Boston, MA, March 9, 1989.

"Results from a Survey of Electronic Board Testing Methods," Digital Equipment Corporation, Andover, MA, September 18, 1990.

"Design Verification and Testing -- A Functional Perspective," Massachusetts Institute of Technology VLSI Seminar, Cambridge, MA, November 20, 1990.

"Design Verification and Testing -- A Functional Perspective," MCC, Austin, TX, December 4, 1990.

"Design Verification and Testing -- A Functional Perspective," Philips Research Laboratories, Eindhoven, The Netherlands, March 1, 1991.

"Design Verification and Testing -- A Functional Perspective," University of Virginia, Charlottesville, VA, July 19, 1991.

"All Tests are not Equally Valuable for Non-Target Defect Detection," Center for Reliable Computing, Stanford University, Stanford, CA, November 13, 1992.

"All Tests are not Equally Valuable for Non-Target Defect Detection," Center for Reliable Computing, Stanford University, Stanford, CA, November 13, 1992.

"New Testing Methods to Enhance Defect Detection using Existing Fault Models and CAD Tools," Computer Engineering Seminar, University of Illinois at Urbana-Champaign, April 15, 1997.

"The Beginning of the End for Stuck-at-Fault Based Testing," Computer-Aided Design Seminar, University of California at Berkeley, October 23, 1997.

"A New Model for Defective Part Level Estimation and its Impact on Automatic Test Pattern Generation," Texas Instruments, Dallas, TX, January 30, 1998.

TUTORIALS:

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), 27th Design Automation Conference Tutorial, Orlando, FL, June 28, 1990.

"Logic Testing and Design for Testability," Rockwell Testing Conference, Newport Beach, CA, January 17, 1991.

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), European Test Conference Tutorial, Munich, Germany, April 10-12, 1991.

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), 5th Annual European Computer Conference Tutorial, Bologna, Italy, May 13-16, 1991.

"Introduction to Integrated Circuit Design" Motorola (William Cannon Site), Austin, Texas, December 3, 5, 9, and 10 1991.

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), Nordic Workshop on Design Verification and Test, Roros, Norway, March 11, 1992.

"Introduction to Integrated Circuit Design" Motorola (Ed Bluestein Site), Austin, Texas, April 28-May 1, 1992.

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), 29th Design Automation Conference Tutorial, Anaheim, CA, June 12, 1992.

"Introduction to Integrated Circuit Design" Motorola (Ed Bluestein Site), Austin, Texas, September 1, 3, 8, and 10 1992.

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), International Test Conference Tutorial, Baltimore, MD, September 20, 1992.

"Introduction to Integrated Circuit Design" Motorola (William Cannon Site), Austin, Texas, December 14-17, 1992.

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), International Test Conference Tutorial, Baltimore, MD, September 20, 1992.

"Testing Digital Circuits and Design for Test," (with T. W. Williams), IEEE International ASIC Conference Tutorial, Rochester, NY, September 28, 1993.

"Techniques for Designing More Testable Logic Networks," (with T. W. Williams), International Test Conference Tutorial, Baltimore, MD, October 17, 1993.

"Introduction to Integrated Circuit Design" Motorola (Ed Bluestein Site), Austin, Texas, September 21-24, 1993.

"Introduction to Integrated Circuit Design" Motorola (Ed Bluestein Site), Austin, Texas, November 29 - December 2, 1993.

"Introduction to Integrated Circuit Design" Motorola (Ed Bluestein Site), Austin, Texas, March 21 - March 24, 1994.

"Introduction to Integrated Circuit Design" Motorola (William Cannon Site), Austin, Texas, April 5 - April 8, 1994.

"Introduction to Integrated Circuit Design" Motorola (Ed Bluestein Site), Austin, Texas, June 28 - July 1, 1994.

"Introduction to Integrated Circuit Design" Motorola (Ed Bluestein Site), Austin, Texas, September 13-16, 1994.

"Testing Digital Circuits and Design Using Scan and Self-Test," (with T. W. Williams), IEEE International ASIC Conference Tutorial, Rochester, NY, September 19-23, 1994.

PATENTS:

"Scan Testable Integrated Circuit" (with V. D. Agrawal), Patent 4,493,077, United States Patent and Trademark Office, issued January 8, 1985.

"Universally Testable Logic Elements and Method for Structural Testing of Logic Circuits Formed of Such Logic Elements," Patent 4,625,310, United States Patent and Trademark Office, issued November 25, 1986.

GRANTS AND CONTRACTS:

University Research Institute, "A New Automatic Test Generation Algorithm," April-August, 1983, $2,136.

Hewlett-Packard Equipment Grant, 1984, $2,800.

Bureau of Engineering Research, "Generalized Graph Operations for CAD Systems," 1984, $3,000.

Microelectronics and Computer Technology Corporation, "Rule Based Automatic Test Pattern Generation Using Boolean Difference Concepts," January 1 - December 31, 1985, $48,976.

GE Calma, "Software License for the TEGAS Logic Simulator," May, 1983 - August, 1986, $135,000 commercial value.

AT&T Information Systems, "Automatic Testing for Faults in Digital Systems," January 1, 1985 - August 31, 1986, $25,000.

Microelectronics and Computer Technology Corporation, "Test Generation for Faults," January 1 - December 31, 1986, $25,000.

International Test Foundation, "Automatic Test Pattern Generation for Delay Faults in Digital Logic Circuits," September 1, 1986 - August 31, 1987, $14,592.

AT&T Information Systems, "Fault Detection in Digital Systems", May 1, 1986 - December 31, 1987, $30,000.

Microelectronics and Computer Technology Corporation, "Continuation of Testing Research," January 1 - December 31, 1987, $25,000.

AT&T Information Systems, "Continuation of Fault Detection Research," May 1, 1987 - December 31, 1988, $25,000.

International Test Foundation, "Test Technology in the Electrical Engineering Curriculum," January 1, 1988 - July 1, 1989, $30,985. ($10,153 for NSF Matching)

Microelectronics and Computer Technology Corporation, "Continuation of Testing Research," January 1 - December 31, 1988, $25,000.

Microelectronics and Computer Technology Corporation, "Testable System Design of Digital Systems and Knowledge Based Structures," (with Xi-an Zhu) April 1, 1988 - March 31, 1989, $92,750.

Office of Naval Research, "Fault-Tolerant Design Techniques for Advanced Digital Architec tures" (with M. Malek), Contract #N00014-86-K-0554, July 1, 1986 - December 31, 1988, $180,000.

National Science Foundation Presidential Young Investigator Award, Grant #MIPS-8552537, June 1, 1986 - May 31, 1991, up to $500,000 (with matching industrial funds).

AT&T Information Systems, "Topological Testing," September 1, 1988 - August 31, 1989, $25,000.

Semiconductor Research Corporation, "The Design of Testable Systems" (with J. Abraham, J. Rahmeh and W. Rogers), SRC Contract, September 1, 1988 - August 31, 1989, $20,000.

Office of Naval Research, "Testing and Fault-Tolerant Design Techniques for Advanced Digital Architectures" (with M. Malek), Contract #N00014-86-K-0554, January 1, 1989 - December 31, 1991, $240,000.

Semiconductor Research Corporation, "The Design of Testable Systems" (with J. Abraham, J. Rahmeh and W. Rogers), SRC Contract #88-DJ-142, January 15, 1989 - January 14, 1990, $250,000.

IBM Corp., "Electronic Testing -- Department Grant" (with J. Abraham), August 1, 1989 - July 31, 1992, $75,000.

Cimflex Teknowledge, "Knowledge Based Design for Testability" (with Xi-an Zhu), June 13 - December 31, 1989, $28,015.

Microelectronics and Computer Technology Corporation, "Continuation of Testing Research," December 7, 1989 - December 31, 1990, $5,000.

Semiconductor Research Corporation, "The Design of Testable Systems" (with J. Abraham, J. Rahmeh, W. Rogers, and L. Pillage), SRC Contract #90-DP-142, January 15, 1990 - January 14, 1991, $350,000.

Texas Advanced Technology Program, "Refined Models of Integrated Circuit Defects Inducing Additional Delays," (with Lawrence T. Pillage), April 24, 1990 - August 31, 1991, $ 137,922.

Semiconductor Research Corporation, "The Design of Testable Systems" (with J. Abraham, J. Rahmeh, W. Rogers, and L. Pillage), SRC Contract #91-DP-142, January 15, 1991 - August 31, 1992, $585,000.

Motorola, Inc., "Timing Analysis for Integrated Circuits," (with Lawrence T. Pillage), November 29, 1991 - December 31, 1992, $10,000.

Office of Naval Research, "Enhanced Timing Analysis for Reliable Wafer Scale Integrated Systems," Contract #N00014-92-J-1723, May 1, 1992 - April 30, 1995, $300,000.

Semiconductor Research Corporation, "The Design of Testable Systems" (with J. Abraham, W. Rogers, and L. Pillage), SRC Contract #92-DP-142, September 1, 1992 - August 31, 1993, $370,000.

Advanced Research Projects Agency (ARPA), "A Unified CAD Tool for Integrated Systems," (with Dean Neikirk and Lawrence T. Pillage) DAAL01-93-K-3317, February 26, 1993 - February 26, 1995, $648,069.

Semiconductor Research Corporation, "The Design of Testable Systems" (with J. Abraham),
SRC Contract #93-DP-142, September 1, 1992 - August 31, 1993, $270,000.

National Science Foundation, "ARI: Development of a Novel Systems Software for Multimedia and High-Performance Computing," Reddy, Mercer, Lu, Cantrell, and Choi, September 15, 1996 - August 31, 1999, $127,450, (Equipment grant).

Texas Advanced Technology Program, "Defect-Directed Test Pattern Generation for Manufacture Testing of Integrated Circuits," January 1, 1998 - December 31, 1999, $139,491.

Semiconductor Research Corporation Custom Research Proposal Sponsored by Texas Instruments, "Automatic Test Pattern Generation for Defect-Directed At-Speed Testing," (with Mike Grimaila) August 1, 2000 - July 31, 2003, $165,000.

U. S. Department of Education, "Meeting the Purposes of Authorizing Statue," (with N. Reddy and K. Watson) August 1, 2001 - July 31, 2004, $ 327,600.

IBM Faculty Partnership Award, "Novel Techniques for Quantifying Confidence during Multi-Processor Verification, Validation, Debug, and Diagnosis," August 1, 2001 - July 31, 2002, $ 25,000.

AMD Research Support Grant, "Integrated Circuit Testing," November 1, 2001 – October 31, 2002, $ 8,000.

Texas Advanced Technology Program, "Integrating Design Verification Techniques with Defect-Oriented ATPG for Very Deep Submicron Systems," April 18, 2002 - December 31, 2003, $139,720.

IBM Faculty Partnership Award, "A New Approach During Multi-Processor Verification and Validation for Estimating Design Correctness," August 1, 2002 - July 31, 2003, $ 25,000.

IBM Faculty Partnership Award, "Quantifying Design Correctness during Multi-Processor Verification and Validation," August 1, 2003 - July 31, 2004, $ 25,000.

IBM Faculty Partnership Award, "A Study of AC Timing Defects: Test Pattern Quality and its Relationship to Real-Time System Errors," August 1, 2004 - July 31, 2005, $ 25,000.

COURSES AT TEXAS A&M UNIVERSITY:

Semester		Course #	Course Title
Fall	2004	EE 680	Testing and Diagnosis of Digital Systems
Fall	2003	EE 652	Switching Theory
Spring	2003	EE 248	Introduction to Digital Logic Design
Spring	2002	EE 248	Introduction to Digital Logic Design
Spring	2001	EE 248	Introduction to Digital Logic Design
Fall	2000	EE 652	Switching Theory
Spring	2000	EE 248	Introduction to Digital Logic Design
Fall	1999	EE 680	Testing and Diagnosis of Digital Systems
Spring	1999	EE 248	Introduction to Digital Logic Design
Fall	1998	EE 652	Switching Theory
Spring	1998	EE 248H	Introduction to Digital Logic Design -- Honors
Fall	1997	EE 680	Testing and Diagnosis of Digital Systems
Spring	1997	EE 248	Introduction to Digital Logic Design
Fall	1996	EE 652	Digital Systems Design
Spring	1996	EE 248	Introduction to Digital Logic Design
Fall	1995	EE 680	Testing and Diagnosis of Digital Systems

COURSES AT THE UNIVERSITY OF TEXAS AT AUSTIN:

Semester		Course #	Course Title
Spring	1995	EE 382M	Topics in Design Verification and Testing
Spring	1994	EE 360M	Digital Systems Engineering II
Fall	1993	EE 382L	Switching Theory
Spring	1993	EE 360M	Digital Systems Engineering II
Fall	1992	EE 382M	Fault Tolerant Computing I
Spring	1992	EE 360M	Digital Systems Engineering II
Fall	1991	EE 382L	Switching Theory
Spring	1991	EE 360M	Digital Systems Engineering II
Fall	1990	EE 382L	Switching Theory
Spring	1990	EE 382M	Fault Tolerant Computing I
Fall	1989	EE 382L	Switching Theory
Spring	1989	EE 382M	Fault Tolerant Computing I
Fall	1988	EE 382L	Switching Theory
Spring	1988	EE 382M	Fault Tolerant Computing I
Fall	1987	EE 382L	Switching Theory
Fall	1987	EE 360M	Digital Systems Engineering II
Spring	1987	EE 382M	Fault Tolerant Computing I
Spring	1986	EE 382M	Fault Tolerant Computing I
Spring	1986	EE 360M	Digital Systems Engineering II
Fall	1985	EE 382L	Switching Theory
Spring	1985	EE 382M	Fault Tolerant Computing I
Fall	1984	EE 382L	Switching Theory
Fall	1984	EE 360M	Digital Systems Engineering II
Spring	1984	EE 382M	Fault Tolerant Computing I
Spring	1984	EE 382M	Fault Tolerant Computing I
Fall	1983	EE 382L	Switching Theory
Summer	1983	EE 382L	Computer Logic Simulation
Spring	1983	EE 360M	Digital Systems Engineering II

PH.D. SUPERVISIONS COMPLETED:

Thomas E. Kirkland		1986	University of Texas at Austin
Ki Soo Hwang		1986	University of Texas at Austin
Rhonda Gaede		1988	University of Texas at Austin
C. T. Glover		1989	University of Texas at Austin
Eun Sei Park		1989	University of Texas at Austin
Kenneth Butler		1990	University of Texas at Austin
Don Ross		1990	University of Texas at Austin
Bret Stewart		1990	University of Texas at Austin
Rohit Kapur		1992	University of Texas at Austin
Mark Heap (with W. A. Rogers)		1993	University of Texas at Austin
Ronn Brashear		1994	University of Texas at Austin
Jaehong Park		1995	University of Texas at Austin
Chanhee Oh		1995	University of Texas at Austin
Li-Chung Wang		1996	University of Texas at Austin
Steve Smith		1996	University of Texas at Austin
Mike Grimaila	August	1999	Texas A&M University

Maximizing Non-Target Defect Detection Using Conventional Stuck-at Fault-Based Automated Test Pattern Generation Tools

Sooryong Lee	August	2003	Texas A&M University

A New ATPG Algorithm to Generate a Compact Test Sets Which Detect Static and Dynamic Defects in VLSI Circuits

Jennifer Dworak	May	2004	Texas A&M University

Modeling Defective Part Level Due to Static and Dynamic Defects Based upon Site Observation and Excitation Balance

M.S. SUPERVISIONS COMPLETED:

Dong Whoan Kim		1984	University of Texas at Austin
Hosung Kim		1984	University of Texas at Austin
Eric J. Schell		1984	University of Texas at Austin
James McKenzie		1985	University of Texas at Austin
Rhonda Gaede		1986	University of Texas at Austin
Ken Butler		1987	University of Texas at Austin
Steve McMahan		1987	University of Texas at Austin
Yi-Feng Lin (Report)		1988	University of Texas at Austin
Tarak M. Parikh		1988	University of Texas at Austin
Wilburn Underwood		1988	University of Texas at Austin
Chih-Teng Hung		1989	University of Texas at Austin
Marvin Denman		1990	University of Texas at Austin
Mark Naiver		1993	University of Texas at Austin
David Carlson		1994	University of Texas at Austin
Mehler, Ronald W.	September	1998	Texas A&M University
Koh, T-Pinn Ronnie	December	1998	Texas A&M University
Jennifer Dworak	May	2000	Texas A&M University

(Texas A&M University Honors Program 1997-1998)

Jason Wicker	December	2001	Texas A&M University

An Analysis of Test Effectiveness via Surrogate Simulation of a Commercial IC

Michael Trinka	August	2003	Texas A&M University

Defect Site Prediction Based Upon Statistical Analysis of Fault Signatures

Bradley Douglas Cobb	December	2003	Texas A&M University

(Texas A&M University Honors Program 2000-2001)

Ordered Partial Decision Diagrams and their use in Manufacture-Test Generation

David Dorsey	December	2003	Texas A&M University

(Texas A&M University Honors Program 2001-2002)

Estimating the Expected Latency to Failure Due to Manufacturing Defects

James Wingfield	December	2003	Texas A&M University

Approaches to Test Set Generation using Binary Decision Diagrams

TEXAS A&M UNIVERSITY UNDERGRADUATE HONORS RESEARCH PROJECT SPONSOR:

Jennifer Dworak	1998	Best Student Presentation Award
Brad Cobb	2000	Best Student Presentation Award
David Dorsey	2002	Best Paper Award
Jeff Cobb	2003	
Jason Vanfickell	2003	
Nate Davis	2004	

PREVIOUS UNDERGRADUATE RESEARCH STUDENTS:

Theresa Huth	Spring 2003 (Winner of an Undergraduate Student Research Best Oral Presentation Award)
Adam Skelton	(Winner of the Thomas S. Gathright Academic Excellence Award Top Junior in the College of Engineering for 2000 - 2001)
John Lee	Fall 2003 & Spring 2004
Cynthia McReynolds	Summer 2003 & Fall 2003
William Charles Price	Fall 2003 & Spring 2004
Justin Ray	Summer 2003 & Fall 2003
Jason Vanfickell	Summer 2003, Fall 2003 & Spring 2004
Jeff Cobb	Fall 2003 & Spring 2004 & Fall 2004
Nate Davis	Fall 2003 & Spring 2004

POST DOCTORAL RESEARCHERS:

Xian Zhu	1988 -- 1990	
Kenneth Butler	1990	(Fall)
Don Ross	1990 -- 1991	
Mike Grimaila	1999 – 2001	
Jennifer Dworak	2004	(Summer and Fall)

SHORT BIOGRAPHY:

M. Ray Mercer is a Professor Emeritus of Electrical and Computer Engineering at Texas A & M University. In September of 2005 he retired as Professor of Electrical and Computer Engineering, Leader of the Computer Engineering Group, and holder of the Computer Engineering Chair. His research interests are centered in computer engineering and include: the computer-aided design of digital systems, design verification, simulation, design for testability, the modeling of logic networks, automatic test pattern generation, distributed computation, communications, and fault-tolerant computing.

Previously, Dr. Mercer worked at: The University of Texas, Austin, TX; AT&T Bell Laboratories, Murray Hill, NJ; Hewlett-Packard Laboratories, Palo Alto, CA; and General Telephone and Electronics, Mountain View, CA. He holds a B.S.E.E. from Texas Tech University, an M.S.E.E. from Stanford University, and a Ph.D. in Electrical Engineering from The University of Texas at Austin. He was the Program Chairman for the 1989 International Test Conference and holds two patents in design for testability. Mercer became a National Science Foundation Presidential Young Investigator in 1986; he has won Best Paper Awards at the International Test Conference (in 1982), the Design Automation Conference (in 1991), and the VLSI Test Symposium (in 1999); he is a Fellow of the Institute of Electrical and Electronics Engineers.

110713