N92-14219

FAULT-TOLERANCE TECHNIQUES FOR HIGH-SPEED FIBER-OPTIC NETWORKS

John DeRuiter*
Honeywell Inc.
Glendale, Arizona

ABSTRACT

Four fiber-optic network topologies (linear bus, ring, central star, and distributed star) are discussed relative to their application to high data throughput, fault-tolerant networks. The topologies are also examined in terms of redundancy and the need to provide for single-point, failure-free (or better) system operation.

Linear bus topology, although traditionally the method of choice for wire systems, presents implementation problems when larger fiber-optic systems are considered. Ring topology works well for high-speed systems when coupled with a token-passing protocol, but it requires a significant increase in protocol complexity to manage system reconfiguration due to ring and node failures. Star topologies offer a natural fault tolerance, without added protocol complexity, while still providing high data throughput capability.

INTRODUCTION

Traditionally, networks for the commercial market have been designed to provide fault tolerance. This fault tolerance, however, has only been provided to a limited extent. That is, a single fault cannot interrupt communications to all nodes but may be allowed to cause the interruption of communications to a single node or a group of nodes. This is less than desirable for aircraft and space applications where there may be critical communications between individual nodes, requiring the total system, not just the network, to be free of single-point failures. Some applications require greater than single-point failure tolerance. For example, the Space Station is required to be operational after two faults. It is therefore desirable that a network design use modular fault-tolerant techniques that can be expanded to greater levels of fault tolerance. As opposed to a commercial office environment, high-speed aerospace applications may require very rapid fault recovery to avoid data loss or excessive delays. Ideally, a network should support autonomous fault-recovery with the fault recovery mechanisms distributed at the individual nodes to provide as rapid a recovery as possible and to avoid centralized system vulnerability.

Many network architectures have been created, based on various fiber-optic-compatible topologies for both commercial and aerospace applications. Commercial systems are characterized by long runs, a relatively large number of nodes, low cost, and limited fault tolerance; aerospace systems, however, are characterized by short runs, a smaller number of nodes, low power, high reliability, and more extensive fault tolerance. This paper examines various fiber-optic topologies, their protocols relative to fault tolerance, and their applicability to the aerospace environment. First, the linear bus topology is discussed; then the ring architecture is examined. Finally, star topologies are addressed.

LINEAR BUS TOPOLOGY

Traditionally, linear bus has been the topology of choice for wire systems. Because of implementation issues, however, it has only limited applicability to fiber-optic networks. In general, the number of nodes that can be supported by a linear bus, without repeaters, is severely limited due to cascaded optical coupler/connector losses and receiver dynamic range and sensitivity limitations.

Since most optical tee couplers are unidirectional devices in which splitting ratios are not reciprocal, a linear bus topology is usually configured with separate couplers and fiber for transmit and receive,[1] as shown in Figure 1. It can be seen that if all transmitters have the same power output and all couplers have the same splitting ratio, the dynamic range requirements imposed on the receiver will increase as the number of nodes in the network increases. In addition, because of the cumulative effect on attenuation of cascading couplers and connectors, receiver sensitivity requirements also increase in proportion to the number of nodes. When considering LED transmitter power, coupler/connector loss, and dynamic range/sensitivity characteristics of available pin

*Staff Engineer; Honeywell Inc., Satellite Systems Operation, Glendale, AZ

diode receivers at high data rates, network size is limited to approximately six nodes if coupler splitting ratios and transmitter powers are fixed. Varying the transmitter power or coupler splitting ratios decreases the dynamic range requirements on the receiver but does not substantially decrease the receiver sensitivity requirements. Accumulated connector, fiber, and coupler losses prevent the linear bus from supporting much greater than seven nodes, even when techniques are used that limit the dynamic range requirements.

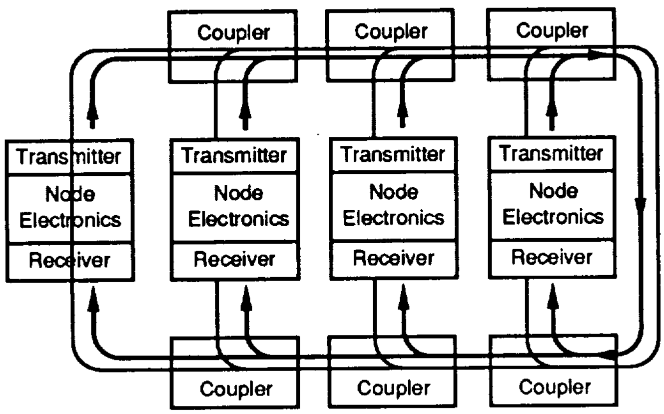

Figure 1. Linear Bus Topology

Linear Bus Fault Tolerance

Fault tolerance for the linear bus requires a duplication of both fiber and couplers. Figure 1 shows that if only a single fiber fails between either a receiver or transmitter and a coupler, only a single node will be affected. If, however, a coupler fails, total network failure can result. It is, therefore, necessary to duplicate all couplers, fibers, and node electronics (as shown in Figure 2) to provide a single fault-tolerant network. This technique can be extended if greater fault tolerance is desired. Should a failure occur in the primary network, however, all activity must be switched to the backup network leaving serviceable node electronics idle. This can be remedied by cross-strapping the primary and

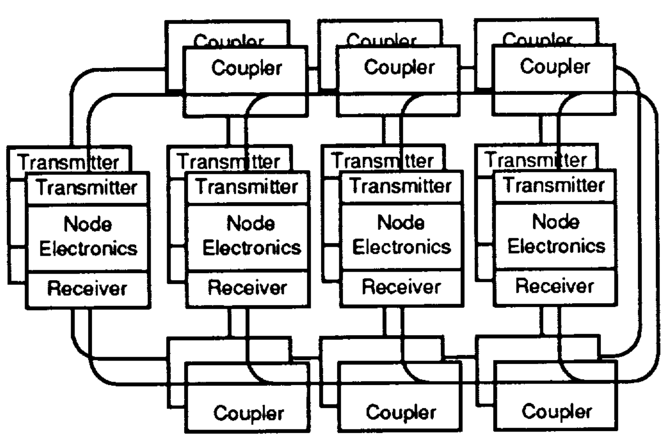

Figure 2. Redundant Linear Bus Topology

backup node electronics to both primary and backup networks. Unfortunately, this adds attenuation to the linear bus, further decreasing the already limited number of nodes that can be supported. The linear bus appears to be less than ideal when applied to any reasonable-size, fiber-optic network with fault-tolerant requirements.

RING TOPOLOGY

Basic ring topology, shown in Figure 3, has the advantage of being a group of point-to-point links, with each node being an active repeater; thus it requires no optical couplers. The dynamic range and sensitivity problems that limit the number of nodes in a linear bus topology are, therefore, substantially eliminated with the ring. Unfortunately, the network is now subject to total failure if any single node or fiber fails. Redundant components can be used to overcome this problem. Interruption of network communication is now, however, a function of active components (repeaters) as opposed to passive components (optical couplers).

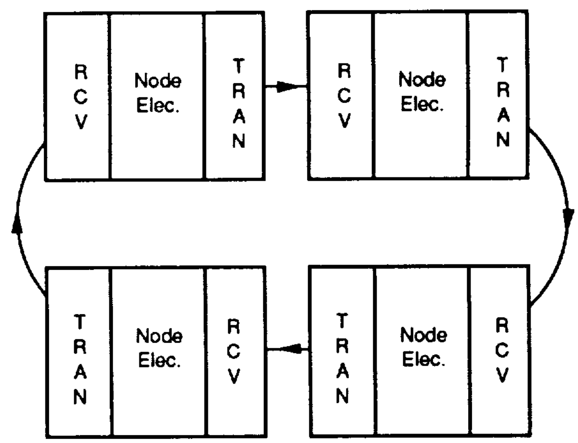

Figure 3. Basic Ring Topology

Ring Protocol

To provide for deterministic operation and high efficiency at high data rates, a token-passing protocol is typically used on a ring. The token ring protocol is based on the idea of a free token circulating around the ring. When a node desires to transmit, it captures the token and then transmits its data. Upon completion of its transmission, the token is reissued. Subsequent stations on the ring then have the opportunity to capture the token and to transmit their own data. Additionally, these token protocols incorporate features to recover from errors on the ring that cause total disruption of network communication (in particular, lost tokens due to bit error rate effects). This is done, however, at the expense of protocol complexity and ring down time. For example, an FDDI system, upon detection of a lost token, requires that all nodes enter a "claim token" mode and, in concert, determine which node has the highest priority and, therefore, the right to transmit and issue a new token. This increases protocol complexity, and, due to the needed cooperation between all nodes, necessitates an interruption in communication to all nodes.

Ring Fault Tolerance

To provide fault tolerance within an optical ring topology, and not just to accommodate soft-error recovery, two additional techniques can be used, including optical bypassing and counter-rotating rings.

A failed node can be bypassed using an optical switch. In a spacecraft application, where power and reliability are critical, it is advantageous to power down any unused nodes both to lower power and to increase reliability. The optical bypass provides a means to circumvent these powered-down nodes. Bypass control can be a completely distributed function, with each node providing autonomous fault detection and bypass. Unfortunately, the optical bypass switch adds attenuation between nodes and, together with optical receiver sensitivity and dynamic range capabilities, limits the number of adjacent nodes that can be bypassed. Only about three adjacent nodes can be bypassed. This is a small number considering a ring's capability of supporting a large number of nodes. In systems where it is desirable to power down a large number of nodes to decrease power consumption and to increase reliability, the ring limits flexibility because care must be taken in how many adjacent nodes are powered down. An additional consideration is that ring operation is interrupted for a finite amount of time because of the bypass switching time. This time can be as great as 25 milliseconds. For high-speed networks, relatively large queues can be required within the node electronics to prevent data loss due to the network communication disruption, which is caused by bypass switching time.

Whereas the optical bypass provides a means for bypassing powered-down or failed nodes, the counter-rotating ring provides for proper ring operation even after a fiber cable has failed. Figure 4 shows how the ring would reconfigure if a cable break should occur. Even though there is a cable break, all nodes can still communicate over a ring that is approximately twice as long as the original. This increases ring latency, but, for aerospace application where run length is relatively short, this effect is insignificant. This provides for single fault-tolerant operation on a system basis if each node is internally dual redundant or if redundant nodes are inserted into the ring. Ring reconfiguration is accomplished by a cooperative effort between all nodes on the ring to locate the break, initiate the necessary reconfiguration, and reinitialize the network. The expense of this cooperative effort, as with recovery from a lost token, is an increase in the network protocol complexity and, as with the bypass, a temporary interruption of network services to all nodes.

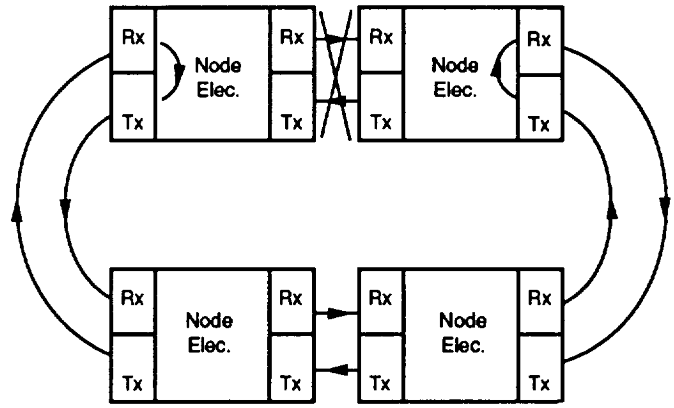

Figure 4. Reconfigurated Counter-rotating Ring (Cable Failure)

The ability of the ring topology to satisfy greater than single fault-tolerant requirements is not a simple extension of the counter-rotating ring technique. It can be solved, however, by adding additional rings. This, like the redundant linear bus, requires that all activity be switched to the backup network. This does not provide the optimum reliability, since it leaves serviceable node electronics idle on the failed ring or rings. Cross-strapping can be implemented to solve this problem, as shown in Figure 5, effectively providing active nodes as a bypass mechanism instead of a simple optical bypass. It is still desirable, however, to incorporate optical bypasses to allow for power down of all node electronics. If a node fails, the network first goes into "claim token" mode, and then into "beacon" mode to identify the failed node. The network manager can then issue a command to the backup node electronics to insert itself into the ring. In this manner, serviceable node electronics are conserved. Unfortunately, the total system is affected by this reconfiguration, not just the failed node, thus incurring an interruption in services to all nodes.

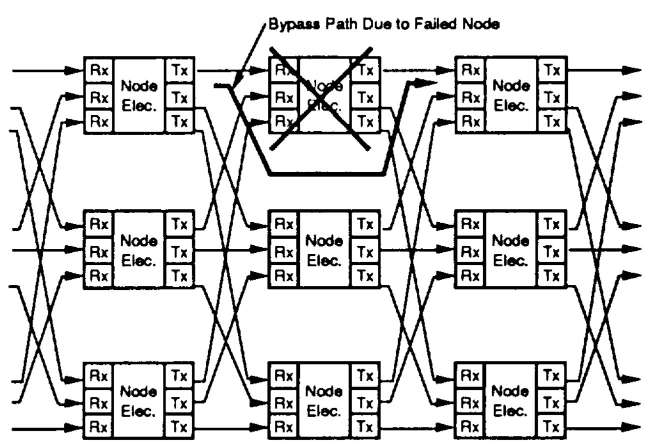

Figure 5. Two-fault-tolerant Cross-strapping for Ring Topology

STAR TOPOLOGIES

Star topologies offer another, and in many cases, better choice for high-speed, fiber-optic, fault-tolerant networks. As shown in Figure 6, a centralized star topology is composed of a variable number of nodes interfaced via a star coupler. This star coupler can be either active or passive. Considering, however, the high reliability requirements of the desired networks, only passive optical star couplers are considered here because of their greater reliability. The star topology, like the ring, overcomes the need for optical receivers with large dynamic ranges. Unlike the ring, however, as the number of required nodes in the network increases, so does the attenuation in the star coupler. This requires greater receiver sensitivity or higher transmitter power for larger networks. Using LED emitters and PIN diode receivers, the star topology can support networks of 50 nodes, which should be quite sufficient for most aerospace applications. Up to 200 nodes can be supported by making use of laser diodes and avalanche photodiodes. Unlike the ring topology, however, the star is not susceptible to total network failure or disruption due to the failure of a single node or fiber. It can also incorporate cross-strapping techniques in its redundant configurations that make more efficient use of system components without added protocol complexity and, therefore, improve both system reliability and fault tolerance. Another advantage of the star topology is that no bypass mechanisms are needed at powered-down nodes, which allows any number of nodes or any sequence of nodes to be powered down. This offers potential power savings and better reliability for those systems that have requirements for a "sleep" mode where a large percentage of the nodes are inactive or not used.

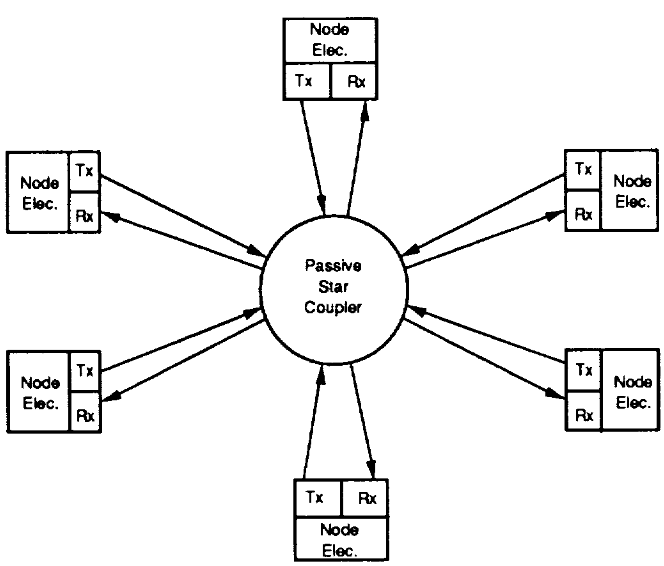

Figure 6. Basic Star Topology

Star Protocol

Both token-passing and contention-type protocols can be implemented on a star topology. At low data rates, 10Mb/s and below, both are efficient.[1] At high data rates, token passing becomes inefficient because of the greater token-passing overhead associated with the star topology. Similarly, at high network loads, traditional contention protocols become inefficient because of the contention resolution algorithms that are used. Two contention-type protocols, however, offer both high efficiency and deterministic operation that make the star topology especially applicable to high-speed, fault-tolerant networks. Both of these protocols (Honeywell's Star*Bus protocol[2,3] and Network Systems' HYPERchannel™) resolve contentions in a deterministic manner via a time-slot cycle. This allows the network to have the efficiency and deterministic properties of a time-slot (virtual token) protocol and the simplicity and fault-tolerant advantages of a contention protocol. That is, since no tokens are passed from node to node, tokens cannot be lost. Fault recovery from lost tokens is, therefore, not necessary; thus, system fault tolerance is enhanced and the protocol and overall system operation are simplified.

Star Fault Tolerance

The star can be made fault tolerant through simple duplication of components. This duplication, as with the redundant linear bus or ring, requires that all activity be switched to the backup network, leaving serviceable node electronics idle. Interruption of network services to all nodes also occurs if a coupler fails due to fault detection and switching time.

Without added protocol complexity, the cross-strapping techniques shown in Figure 7 can be used with the star topology to make more efficient use of system components, while providing virtually instantaneous fault recovery in the event of a transmitter, receiver, fiber, or coupler failure. Node electronics can consist of nonredundant elements, as shown in Figure 8, or internally redundant elements, as shown in Figure 7. Nonredundant elements offer the advantages of providing a building-block approach to fault tolerance and minimal impact to the user that does not require redundancy. The use of nonredundant elements, however, requires additional taps on the star couplers: two per node for a dual system, three per node for a triple, etc. The use of internally redundant elements implies added complexity, but limits the number of taps necessary on the star coupler to one per node and, therefore, has an advantage relative to network loss budget. In both cases, the cross-strap at the optical media works the same. As shown in Figure 9, both transmitters generate identical data, with one transmitter being interfaced to the primary coupler and the second being interfaced to a backup coupler. At the receiving node, both receivers are active, with only one being selected. If both receivers pick up a signal, priority is given to receiver A. If only one signal is present (indicating a failed transmitter, fiber, coupler, or receiver), the active receiver output will be selected. In this manner, with dual transmitters operating in parallel in the sending node and dual receivers selecting the active channel in the receiving node, virtually instantaneous fault recovery is provided. No channel selection is performed at the system level; thus overall system management is simplified. Because of the fault-tolerant nature of this configuration, override capability is provided to allow for test functions.[2]

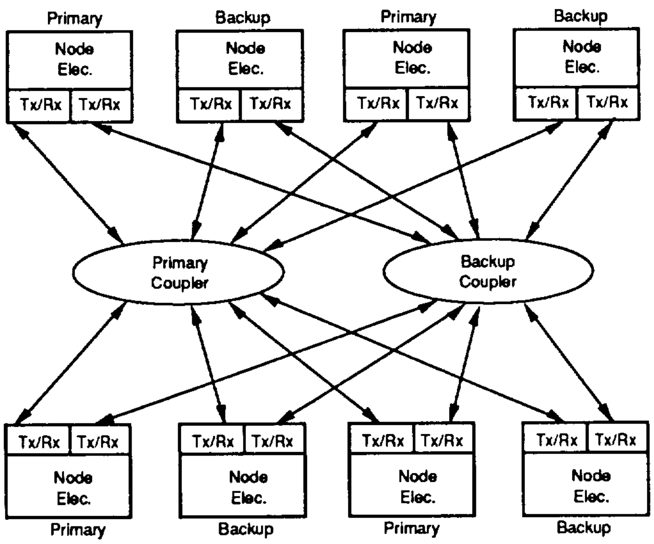

Figure 7. Dual-redundant Cross-strapped Star (With Internally Redundant Node Electronics)

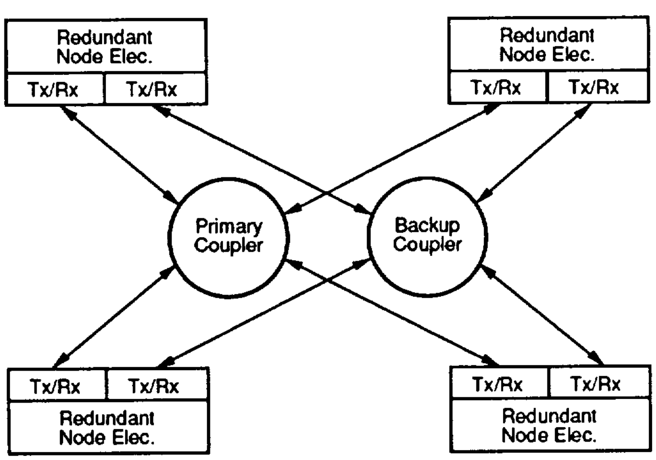

Figure 8. Dual-redundant Cross-strapped Star

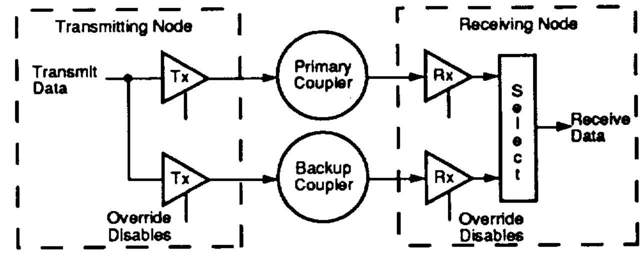

Primary Active	Backup Active	Selected Channel
1	1	Primary
1	0	Primary
0	1	Backup
0	0	Neither

Figure 9. Cross-strap Operation

Another star topology is represented by the distributed star[1] shown in Figure 10. This configuration offers greater fault tolerance than the central star in that no single optical coupler can cause the interruption of communication to all nodes. It does, however, have greater connector and excess coupler losses because of the cascading of couplers. This, in general, limits the number of nodes it can support relative to the central star approach. To achieve total system fault tolerance, the same component duplication and cross-strapping techniques (as previously described for the central star topology) can be used.

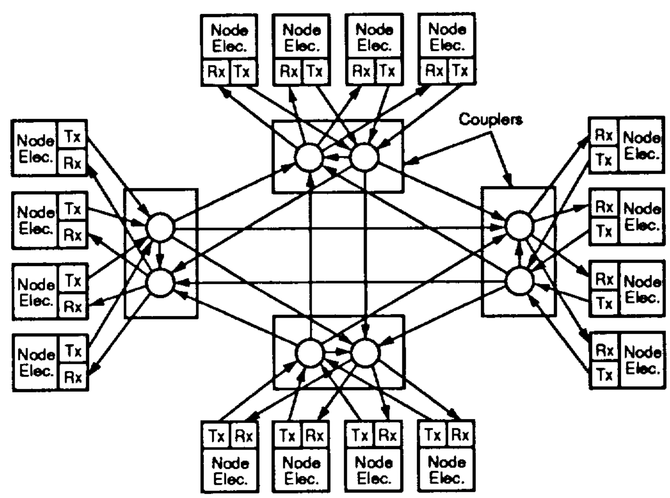

Figure 10. Sixteen-node Distributed Star

AUTONOMOUS NODE FAULT RECOVERY

Some fault-recovery tasks must be performed regardless of topology. One of these tasks is switching from primary to backup node electronics. This can be done on a system basis, for example, as discussed previously for the ring topology. It is, however, more desirable to provide autonomous fault recovery at each node, thus minimizing the functions required of the network manager, not just in the ring but in all topologies.

Figure 11 shows a detailed block diagram of a possible implementation for an autonomous switchover scheme between two nonredundant units. Each unit is identical, with the "primary ID" causing one to power up as the primary and the "backup ID" causing the other to power up as the backup. The primary is fully powered on and, therefore, fully functional. The backup is in "standby", with only its power-up control, toggle and override detectors, and receivers powered up. The backup monitors the primary's health via the toggling health signal between the two units. The CPU in the primary evaluates built-in-test results and pulses the toggle generator if all tests pass. The presence of the toggling signal is therefore the result of proper operation. The lack of a toggle from the primary will cause the primary to go into "standby" and the backup to become active and take over the node. An override command can be received, via the network, to provide for switchover testing and contingency operations in the event that a failure is not detected or detected in error. To ensure that both primary and backup are not powered simultaneously, the power switch and the override detector are made redundant. Also, the primary toggle detector will put the primary into standby if the backup powers up in error and the backup toggle becomes active. The cross-strapped toggles, therefore, provide a flip-flop type of configuration, with the primary and backup always in opposite states.

GENERAL POWER AND RELIABILITY CONSIDERATIONS

A companion issue to fault tolerance and redundancy is reliability. The intended result of providing redundancy is to enhance overall system reliability. As system components become more unreliable, greater levels of redundancy are necessary to maintain overall system reliability. Reliability is, in general, affected not only by component reliability, but also by circuit complexity and power dissipation. As circuit complexity goes up, component count goes up and reliability goes down. As power dissipation increases, component junction temperatures increase and reliability goes down. Basic differences in protocol complexity have already been discussed and are relatively clear cut. Power dissipation differences between topologies and protocols are, however, more subtle and are discussed here.

The basic nature of a ring topology, coupled with a token-passing protocol, implies a greater power usage than a star topology coupled with a broadcast protocol. In a star topology, only one transmitter is active at any one time. When a node wishes to transmit, it simply monitors the network for activity and transmits its frame when the network is free. This transmission is then received, via the star coupler, by all nodes and requires no action from other than

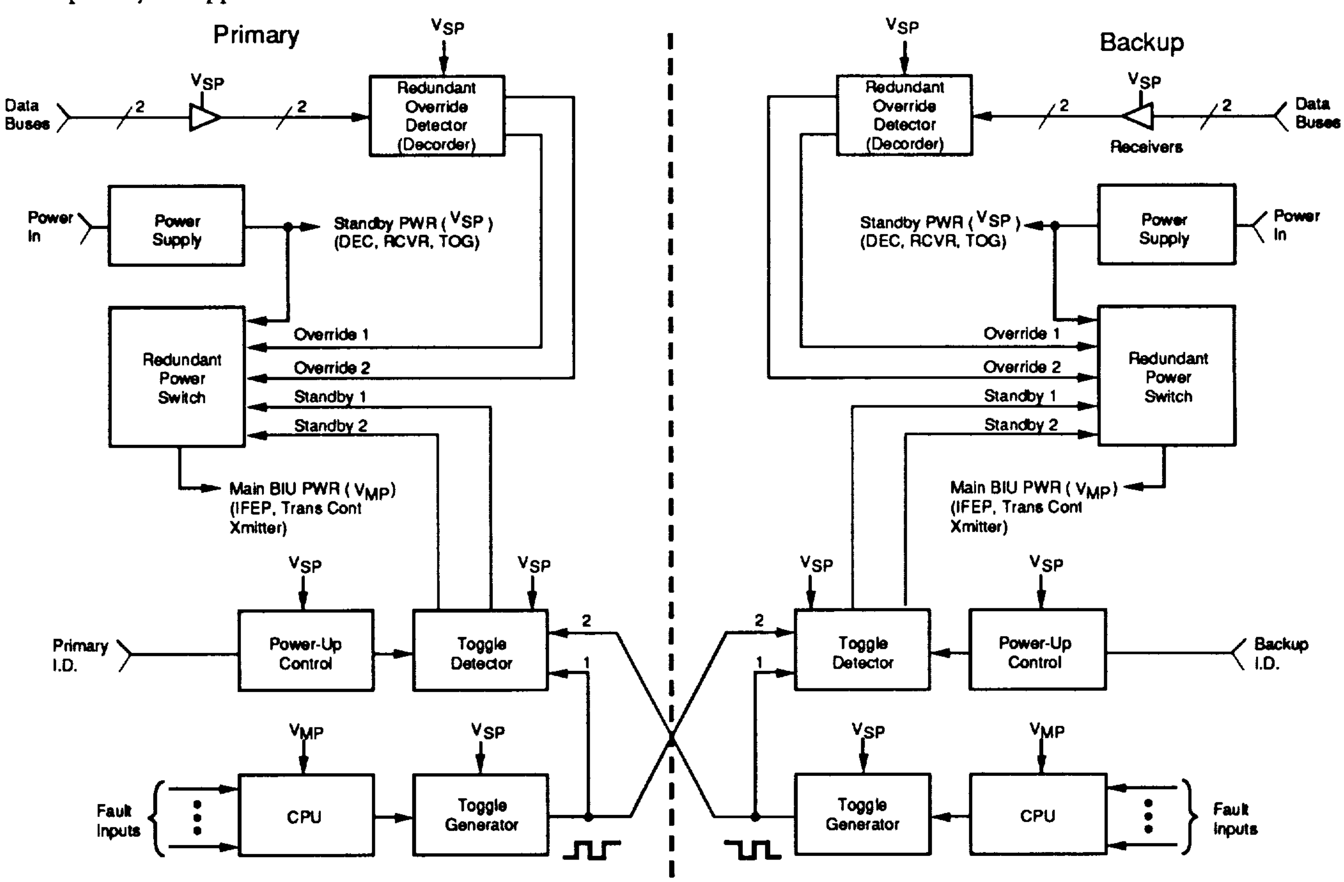

Figure 11. Switchover Implementation

the transmitting and receiving nodes. A ring, on the other hand, requires all nodes to participate in the data transfer. A frame generated by a node in a ring, however, must be repeated by all nodes on the network. This requires a greater duty cycle at each node, with each node having to transmit all frames even though they are not locally originated. In the worst case, should the physical ring be short relative to the transmitted frame, all transmitters will be active simultaneously. For example, at 100 Mb/s, a ring with a one-kilometer circumference requires all transmitters to be active simultaneously when a frame of only 500 bits or greater is circulating on the network. For aerospace applications, where runs are short, this worst-case condition is typical, not merely an exception. Each node in a ring must, therefore, run at a substantially higher duty cycle than each node in a comparable star topology; thus a higher power dissipation is incurred. This causes junction temperatures to elevate and reliability to suffer.

Other power/reliability considerations involve the protocols themselves, without regard to topology. Power-strobing techniques have long been used, in electronics intended for space applications, to reduce power dissipation, and, subsequently, to increase reliability. Circuitry is powered on only when operation is required. This makes power dissipation, and therefore reliability, dependent on duty cycle. The basic nature of a network, in which each node occupies only a portion of the network bandwidth, makes the application of power strobing beneficial. Protocols intended for use in high-reliability systems with only limited power available should be designed to allow the use of these power-strobing techniques while still maintaining high performance. Whether for a linear bus, a ring, or a star topology, the selected protocols should allow for a substantial portion of the circuitry in individual nodes to be powered off when no data transactions are occurring.

SUMMARY

Three basic topologies, relative to their application to fault-tolerant, high-speed networks for aerospace applications, have been examined. Of these three, both the ring and the star are viable candidates. The linear bus presents implementation problems for all but the smallest networks because it is limited in the number of nodes it can support. The ring can support the largest number of nodes and can easily support high data rates and deterministic operation. It can also support various levels of fault tolerance but does so at the expense of fault recovery time and an increase in media access and network management protocol complexity. The star topologies offer a better choice, providing more inherent fault tolerance, while still providing support for high data rates, deterministic operation, and a relatively large network size. The star topology also provides an inherently lower power dissipation; only one node is required to transmit a frame as opposed to the ring where all nodes must repeat the frame. Similarly, since it requires no bypasses for powered-down nodes, the star topology offers potential power savings and better reliability for those systems requiring a "sleep" mode where a significant percentage of the nodes are inactive during a particular mission phase.

REFERENCES

1. Uhlhorn, R.W., "Fiber Optic Architectures and Components for High Data Rate Buses for Military Applications," FOC/LAN 1984, pages 265-272.

2. Varga, Mike, "A Robust 100Mb/s Network for Avionics Applications," NAECON, May 23, 1985.

3. DeRuiter, John, "A Robust Fiber-Optic Network for Space Applications," OE-LAS '86, Vol. 616, January 1986, pages 214-221.