Redundant Ring Technology for Industrial Ethernet Applications

Alvis Chen, Moxa Business Development Manager
alvis.chen@moxa.com

"Redundancy" is now one of the most important aspects of the industrial automation networks used for applications such as power utilities, transportation, and surveillance. In this white paper, we first discuss traditional redundant technology, including their network topologies and limitations. Next, we introduce "redundant ring" technology, and point out the advantages that redundant rings have over traditional systems. In closing, we highlight the key features of industrial Ethernet switches, and show how to select the redundant ring switch that provides your industrial application with the best performance.

Overview of Traditional Redundant Ethernet Technology

Why do Ethernet networks need redundancy?

It goes without saying that industrial Ethernet networks must provide "reliable data transmission." Unfortunately however, the industrial environments that accompany applications such as power stations, automated automobile factories, and trains, are subject to electrical noise, wide temperature variations, vibrations, and all kinds of electrical/physical interference. Since any of these factors can cause communication failures, it is essential that engineers use a "robust" and "redundant" design for their communication network.

Released on November 1, 2007

How to contact Moxa
Tel: 1-714-528-6777
Fax: 1-714-528-6778
Web: www.moxa.com
Email: info@moxa.com

Consider the following example. An energy company wants to create an Ethernet infrastructure for monitoring and controlling three substations from a central control room located at the company's headquarters. The most straightforward topology that can be used for this type of Ethernet infrastructure is the line topology (Figure 1).

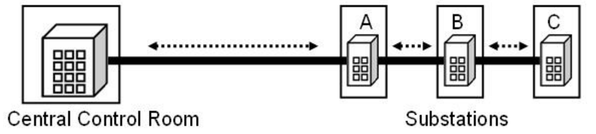

Figure 1: Line topology for substation control

Although line topologies work well in ideal situations, the conditions found in harsh industrial environments can cause the connection to fail. The worst case scenario is if the connection between central control room and substation A fails, since in this case the central control room would not be able to connect to substations B and C, either (Figure 2).

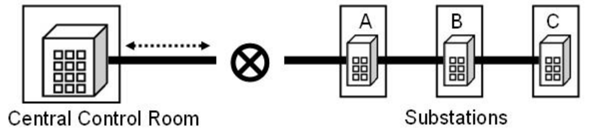

Figure 2: Line topology failure

Mesh-type redundant Ethernet topology

One type of redundant network that can be used to avoid total system failure is called the "mesh-type" redundant topology. Figure 3 shows a complete mesh-type redundant network. Note that in this case, each switch connects to every other switch in the network.

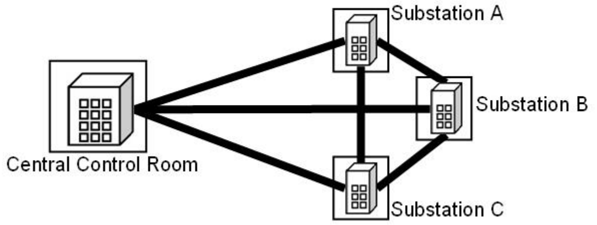

Figure 3: Mesh-type redundant network

In this case, if the connection between the central control room and substation A fails, signals will be transmitted through alternate paths, such as between B and A, or C and A (see the dotted lines in Figure 4). These alternate paths are the so-called "redundant" paths.

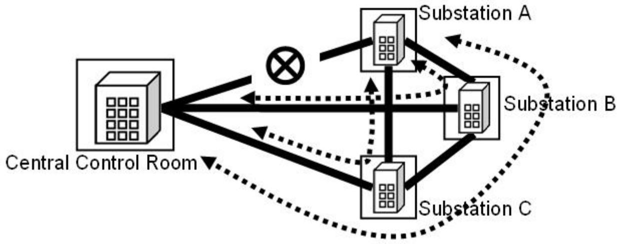

Figure 4: Redundant paths for a mesh-type network

One of the problems with the mesh-type topology is that the Ethernet network could experience "looping," which occurs when signals can travel along more than one path when transmitted from one device to another. For example, as illustrated in Figure 5, a switch in the central control room could connect to the switch in substation A both directly and indirectly. The existence of both paths would cause non-stop data packet looping, and as a consequence decrease the performance of the network.

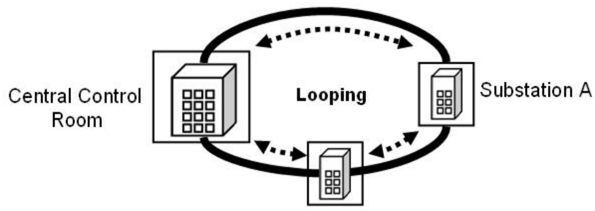

Figure 5: Looping problem in the network

Spanning Tree Protocol

One of the technologies developed to handle Ethernet redundancy is called the IEEE 802.1D Spanning Tree Protocol, or STP for short. STP is used to define certain paths in an Ethernet mesh topology as backup paths. To prevent “looping,” STP sets a single communication path between each switch in the Ethernet network. Communication over the other “redundant” paths is blocked temporarily by STP.

Figure 6 shows an example of a redundant Ethernet network that uses STP. The switches in the network negotiate with each other to decide which paths will be active and which will be designated as redundant. Packets are transmitted over active paths (solid lines), whereas signals are blocked from being transmitted over the redundant paths (dotted lines).

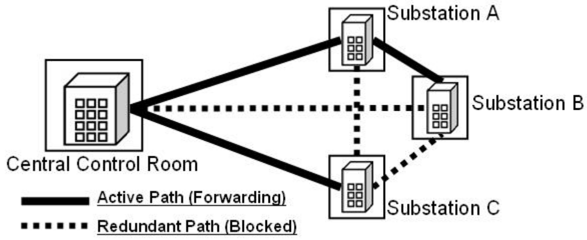

Figure 6: Active and redundant paths in a mesh-type network

When STP is running, packets sent from switch B to the central control room must pass through switch A. even though there is a direct physical connection from switch B to the central control room. As pointed out earlier, STP prevents the “looping” problem posed by the mesh-type system. When one of the paths gets disconnected, or a device in the network fails, STP automatically re-arranges the connections and uses the redundant path. As is illustrated in Figure 7, when the connection from the central control room fails, the path from substation B to the central control room is activated, allowing signals to be transmitted from substation A to the central control room by way of substation B.

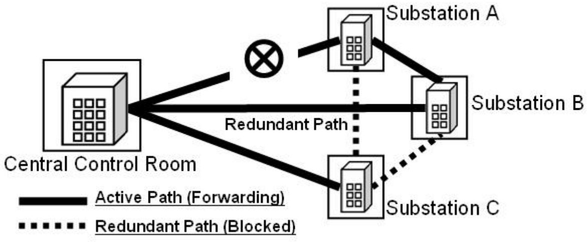

Figure 7: Activation of a redundant path

Limitations of mesh topologies and STP

Although STP fixes the looping problem, when the “mesh topology” and “STP” are used in the real world, the high cost of installing so much wiring can become prohibitive, and the recovery time is typically too long for industrial applications.

(A) Wiring cost:
Figure 8 shows the lengths of wiring required to set up a mesh-type network between the central control room and the three substations (A, B, and C).

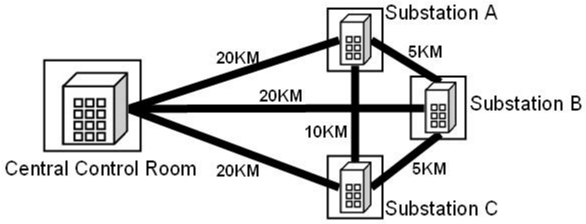

Figure 8: Wiring lengths

In this case, the total amount of wiring needed for the mesh-type network is 20+20+20+10+5+5 = 80 km. However, if the line topology can be used, the amount of wiring required is only 20+5+5 = 30 km (see Figure 1). Another cost to consider is the number of ports needed for the connection. In this case, the “mesh-type” topology uses [3 ports per switch] x [4 switches] = 12 ports, whereas the number of ports needed for a “line-type” topology is 6 (1 port for the control room switch, plus 2 ports for each of switches A and B, plus 1 port for switch C). From this simple analysis, we conclude that the mesh-type topology would cost more then double what the line-type topology costs.

(B) Recovery Time:
The “recovery time” of a redundant network is the time required for the network to activate redundant paths when one of the primary paths becomes inactive. This change occurs when one of the switches is turned on or off, when an active link fails, or an inactive link is plugged in and becomes activate. The standard recovery time for STP is on the order of 30 to 90 seconds. This is simply too long for industrial applications, such as the Ethernet network for large ships, or the production line in an automated factory.

For this reason, IEEE 802.1W Rapid Spanning Tree Protocol (RSTP) was developed. This newer protocol has all the advantages of IEEE 802.1D, but in addition provides higher performance. The average recovery time of an RSTP network can be as much as 1 second. However, for some “mission critical” industrial applications, such as process control and military system, the recovery time must less then 100 milliseconds to ensure the reliability of the Ethernet network.

Ring-type Redundant Ethernet

In order to overcome the limitations of mesh-type topologies in an industrial Ethernet infrastructure, many users have implemented a so-called ring topology. Figure 9 shows how a ring-type redundant Ethernet network would look for the control room plus substations example.

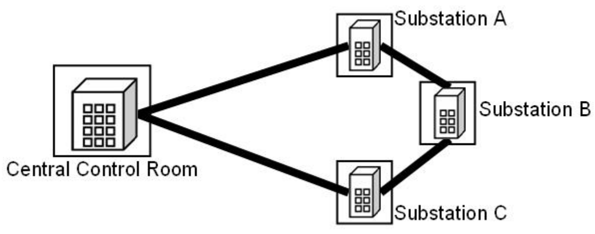

Figure 9: Ring-type network

To avoid looping problems, a ring-type network must implement STP, RSTP, or a proprietary recovery protocol designed exclusively for ring-type networks. For example, we could block data transmission on the segment that connects switches B and C, as shown in Figure 10.

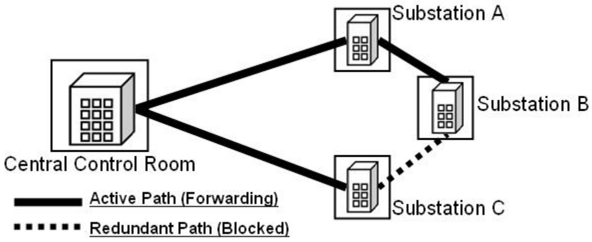

Figure 10: Ring-type network with one segment blocked

Ring topologies have a lower cost for cabling and installation

One of the major benefits of a ring topology is the lower cost of cabling and installation. In Figure 8 we showed that a mesh-type topology for this example would require 80 km of wiring. However, as shown in Figure 11, the ring topology requires only 50 km of wiring. Note that although the line topology requires even less wiring than the ring topology, line topologies do not support redundancy.

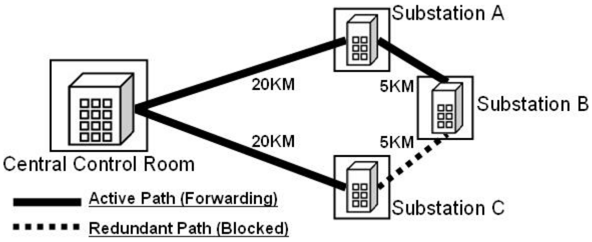

Figure 11: Path lengths for a ring topology

Hierarchical Design

Another benefit of ring topologies is that complex networks can be subdivided into smaller sub-networks. In fact, real systems are usually divided based on location and/or function. Figure 12 shows a typical example of how a modern substation network infrastructure could be subdivided. In this case, the network is made up of three levels: the station level, bay level, and process level.

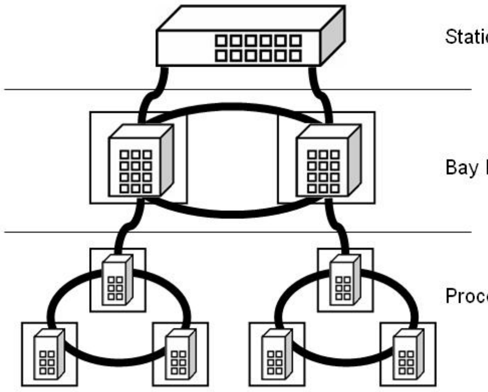

Figure 12: Modern substation network

Ring topologies are ideal for realizing hierarchical structures made up of multiple levels, and different management functions can be used on the different levels. Since most data packets stay within one of the levels, th e hierarchical structure provides better performance as far as Ethernet traffic is considered, and as a result increases the reliability of the system as a whole.

Millisecond high-speed recovery time

Although RSTP provided a significant improvement in recovery time compared to STP, it is still not fast enough for some mission critical applications. For this reason, manufacturers of industrial Ethernet switches have developed proprietary protocols to decrease the recovery time down to the millisecond range.

We point out two ways that can be used to reduce the recovery time. The first method involves improving the performance of RSTP in the manufacturer's own switches. The second method uses proprietary control packets to check the status of the ring. In the following table, we list the proprietary redundant ring protocols available from different manufacturers.

Manufacturer	***Proprietary Redundant Ring Protocol***
Moxa	Turbo Ring
Hirschmann	Hiper Ring
RuggedCom	eRSTP
SIXNET	Real-time Ring

In the following paragraphs we describe six major factors that affect the recovery time of a ring network.

Time required to detect link up/down

The time required to detect link up/down depends mostly on the hardware used in the switches. For example, the transmission medium (copper or fiber) and transmission speed (current options are 10, 100, and 1000 Mbps) both affect the recovery time.

Time required to purge the forwarding table

When a link up/down occurs, the topology changes, or other problems with the infrastructure arise, the forwarding table (sometimes referred to as the Forwarding Database, or FDB) should be changed. However, before the forwarding table is changed, the system must first clear, or purge, the table. The amount of time required to purge the table increases with the size of the table.

Propagation delay

"Propagation delay" refers to the transmission delay that occurs when a packet is transmitted through the network device. It should be obvious that the propagation delay of a network increases as the number of switches increase.

Redundancy for ring-to-ring applications

Because of certain maintenance issues, complex Ethernet networks must be divided into a hierarchical system of redundant ring subsystems. For this reason, how the redundancy protocol handles the connections between sub-rings is important. "Ring Coupling," "Dual-Homing," and "Dual-Ring" are the functions used most often to handle dual ring redundancy.

Ring Coupling

"Ring Coupling" supports separating distributed devices into different smaller redundant rings. Figure 13 shows the ring coupling structure, in which ring A and ring B are connected by a primary path (solid line) and backup path (dotted line). If the primary path fails, the system reverts to using the backup path to connect ring A to ring B.

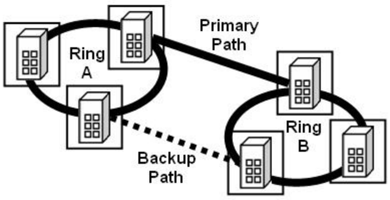

Figure 13: Ring Coupling structure

Dual-Homing

"Dual-Homing" is used for special installation environments. In this case, a single switch in one of the rings connects to two separate switches in the adjacent ring (see Figure 14). The backup path will be activated in the event of a failure of the operating connection (the primary path).

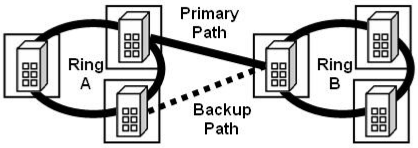

Figure 14: Dual-Homing

Dual-Ring

"Dual-Ring" provides the most cost-effective, and also the most reliable solution. In this case, one switch essentially belongs to both rings (see Figure 15). The dual-ring topology is useful for applications that present cabling difficulties.

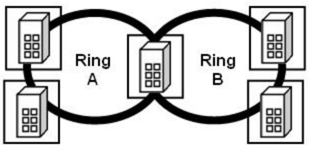

Figure 15: Dual-Ring

The Performance of Ring-type Redundant Ethernet Networks

We have learned that ring-type redundant Ethernet networks provide many benefits compared to traditional topologies. The benefits include faster recovery time, lower installation costs, convenient hierarchical management, and ring-to-ring redundant functions. In this section, we point out three important factors that should be considered when choosing switches that provide the best performance to cost ratio.

Testing report for large numbers of switches connected to one ring

When an Ethernet switch provider announces that the recovery time of their switches is only a few milliseconds, the next question you should ask is this: How many switches can I install in one ring and still get the benefit of such a short recovery time? As we pointed out earlier, the propagation delay of a network increases as the number of switches in the system increases. Therefore, the best way to determine the maximum number of switches that can be installed in one ring is by conducting an actual test that uses a large number of switches (see Figure 15). If such test data is not provided, then the so-called "recovery time" announced by the manufacturer is based on a less reliable theoretical calculation.

Figure 15: Recovery time testing with a large number of switches

Different interfaces and different speed modes

Most published recovery times assume that all switches in the network are configured to use the fastest transmission speed supported by the switches. However, real-world applications can use a combination of interfaces (copper and fiber), and a combination of transmission speeds. Different transmission speeds (10, 100, or 1000 Mbps) can result in different in recovery times. For this reason, be sure to check the specifics of published recovery times to make sure that the numbers apply to your own network.

Testing procedure

The results of a recovery time test are dependent on the testing procedure. When reading test reports, you should pay attention to the following two factors:
(1) The manufacturer should use qualified testing equipment, such as SmartBits, to produce their recovery time results, and the equipment should be calibrated before running the tests.
(2) The forwarding table or MAC address table should be purged while running the test. This is the only way to determine a recovery time that applies even to worst-case scenarios.

Conclusion

High speed redundancy networks have become a key issue when building industrial Ethernet networks. Traditional mesh-type topologies and STP/RSTP redundant protocols do not provide a recovery time short enough for industrial applications, and come with a much higher cost for cabling and installation. Therefore, more and more industrial switch manufacturers now provide high-speed recovery times and redundant ring switches as the best solution for industrial Ethernet infrastructures. However, it is not sufficient to consider only the maximum recovery time. Conducting actual tests with large numbers of switches is the only reliable means of determining the recovery time of redundant networks. For complex networks, the switches you use must support a variety of transmission speeds (10, 100, 1000 Mbps), transmission mediums (copper and fiber), and ring-to-ring redundant topologies (Ring Coupling, Dual-Homing, and Dual-Ring) to set up a reliable and redundant industrial Ethernet network.

Another concern is the demands that are made of industrial automation applications at the management level, control level, and device level. Data loss and time delay are both critical issues for automation design at the control level and device level, but are of particular concern at the device level. For example, on-line processing and metering applications require short time delays for smooth operation, and for this reason, rapid fault recovery with no time delay or data loss is required. Since industrial Ethernet technology penetrates from the control level down to the device level, a faster fault tolerance is needed for industrial Ethernet networks. To meet such a strict requirement, Moxa's new generation of Turbo Ring technology offers a significant reduction in recovery time—20 ms at a 250-switch load—to ensure fast and constant network connections for device-level applications.

