

A Comparison of Dispersed Topologies for Ethernet

Roger Moore
RuggedCom Inc. – Industrial Strength Networks
Concord Ontario Canada

1 Introduction

The trend for automation and control networks in various industries is to use Ethernet as the physical layer. This leverages a vast array of both networking equipment and end devices. However, deploying a real-time control Ethernet network in a geographically dispersed area has a unique set of challenges. When the distance between end stations is many kilometers, the network is still logically a LAN, but physically is anything but local. Cabling costs can be significant in such a network; choosing a topology that minimizes these costs yet provides for a robust, fault tolerant network is more difficult. This paper compares several different topologies that address these needs.

1.1 Goals of the Topology

- Eliminate single point of failure for cabling
- Minimize single point of failure for equipment
- Provide deterministic peer to peer response time
- Minimize peer to peer latency
- Maximize allowable bandwidth per station
- Provide uplink to central management facilities
- Minimize network management
- Minimize cabling cost

1.2 Assumptions and Terminology

The physical topology of the network is obviously dependent on the particular application. For the purposes of this paper a grid will be used; this allows some ballpark estimates of cabling requirements. See figure 1.

The network is a homogeneous, micro-segmented Ethernet network. All long distance links have a 100F_x, single mode fiber optic (1300 nm) physical connection.

Terminology

Bridge: Same as switch.

Field Switch: A 'hardened' Ethernet edge switch. May or may not need to support STP depending upon topology. Requires a minimum of two single mode (1300nm) fiber optic ports.

Backbone: An Ethernet backbone switch with many ports. Must support STP. All connections to field switches must be single mode (1300 nm) fiber optics ports.

End Device: An Ethernet enabled device such as an IED(intelligent electronic device), PLC(programmable logic controller), 2070(traffic controller), Ethernet camera, etc.

Station: A group of end devices connected to the network via a field switch; represented by a numbered circle on figure 1.

nS: Total number of stations

dS: Average distance between stations

dB: Average distance from station to backbone; estimated at $3 \times dS$ for this paper.

A-FS1: Refers to end device A on field switch 1.

Figure 1: Grid Layout of Stations

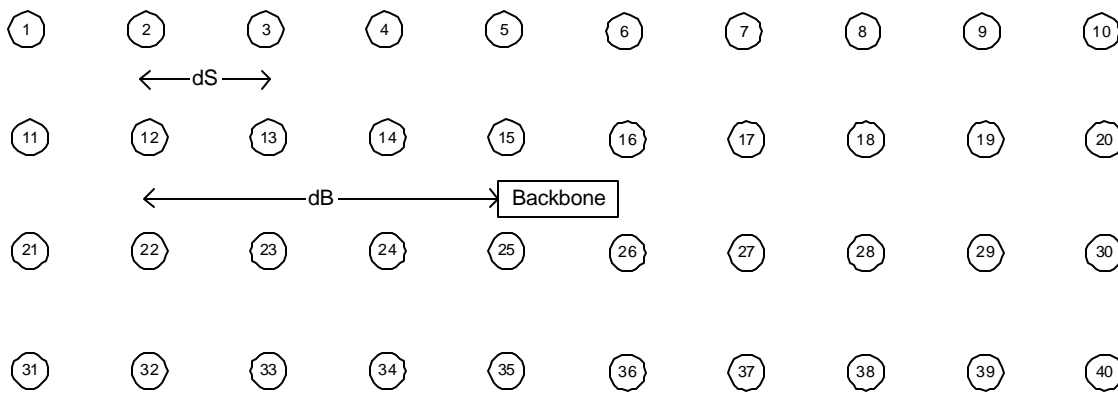
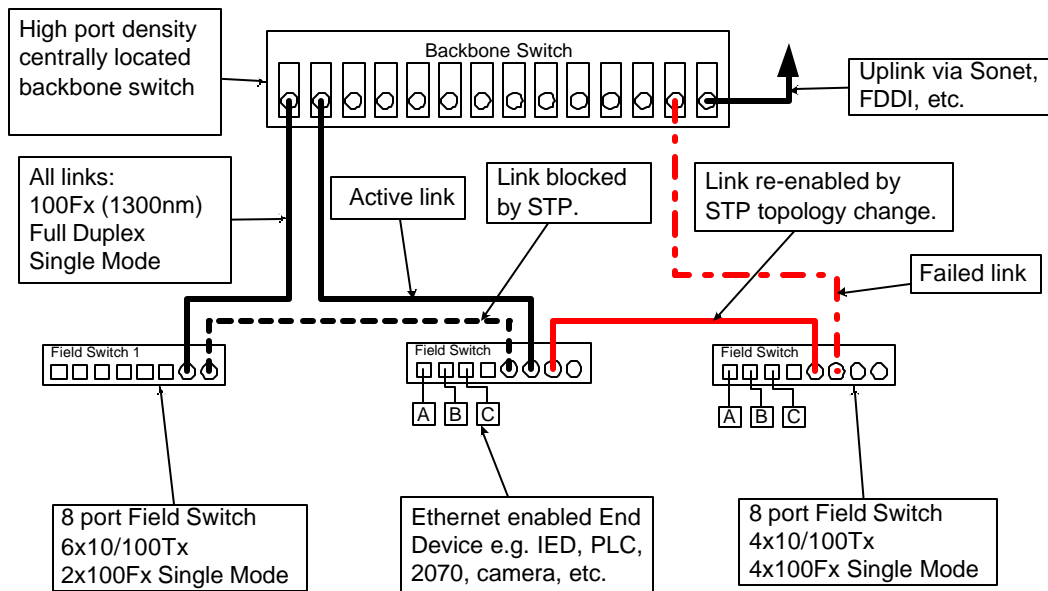


Figure 2: Legend For Topology Diagrams



2 Spanning Tree Protocol Refresher

A bridged Ethernet network cannot have loops. A loop results in what is coined a 'broadcast storm' where frames circulate around the loop forever rendering the network useless. However, loops provide a simple way to have redundant data links thus improving fault tolerance. The Spanning Tree Protocol (STP) allows the network to have physical loops but prevents traffic from circulating on those loops by blocking links.

STP support by a switch implies that the switch is 'managed' and implies higher cost for the device. Managed switches generally come with numerous additional features such as:

- Command line interface for configuring device via RS232 or telnet
- Port configuration
- RMON Ethernet statistics
- SNMP
- VLAN (802.1Q), priority queuing (802.1p)

STP has two primary parameters that must be configured for each switch on the network. They are: *bridge priority*, and *port priority*.

The switch with the lowest bridge priority becomes the *root bridge* in the network. The root bridge is the logical center of the network though not necessarily the physical center. Think of an upside down tree – the root bridge is the base. The root bridge may change over time as switches are removed i.e. fail.

The end goal of STP is to ensure that only one switch is responsible for forwarding traffic from the direction of the root onto any given link. If there is only one active path from the root to a link then there will be no loops in the topology. The switch responsible for forwarding traffic onto a link is called the *designated bridge* for that link.

For a designated bridge there are three types of ports:

A *designated port* is one that forwards traffic from the root onto a link.

A *root port* provides connectivity to the root bridge.

All other ports are inactive or disabled by STP.

STP determines the designated and root ports by determining the lowest cost path (based on link speed) back to the root bridge. Since this paper assumes a homogeneous network all links have equal cost so STP will choose these ports based on the minimum number of 'hops' back to the root.

Note that the presence of loops in the network does not imply that all switches on a loop support STP. As long as at least one switch on any given loop supports STP, the network will work. For the topologies discussed in this paper, STP support is always required from the backbone switch. The need for STP support on the field switches depends on the topology.

3 Small Loops Topology


The small loops topology consists of two stations per loop with a single backbone switch as shown in figure 3. Each station has a direct connection to the backbone and a connection to one neighboring field switch.

This topology does not require STP support from the field switches because the maximum number of hops between any pair of devices is five. The number of hops varies between three and five depending on the devices communicating. For example, between device A on FS1 and device C on FS39 there are three hops. Between device A on FS2 and device C on FS40 there are five hops.

The maximum sustained bandwidth acceptable from a station is 50 Mbps.

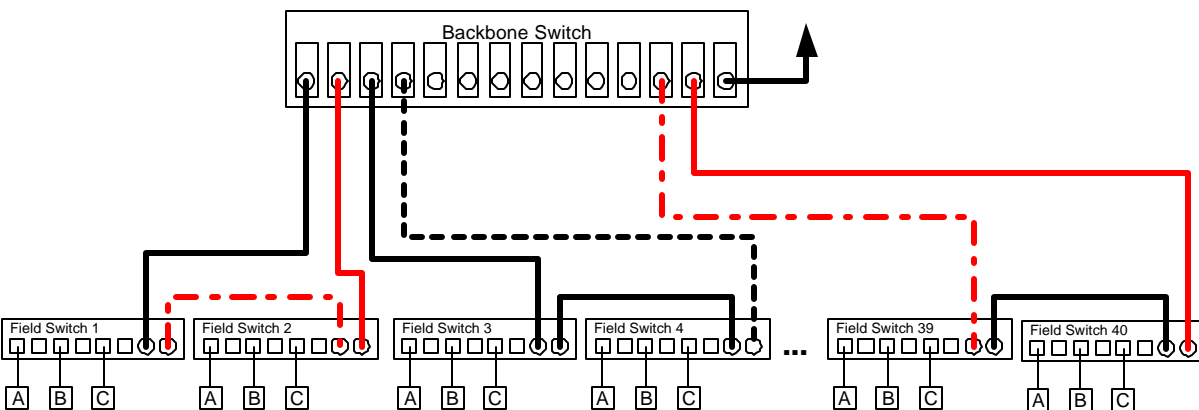
The total cable distance is: $(nS \cdot dB + nS \cdot dS/2) = 3.5 \cdot nS \cdot dS$

Figure 3: Small Loops Topology



A failure of any single link in this topology is recoverable via STP. Failure of two links on a loop will isolate devices on that loop. Failure of the backbone will isolate all loops. Figure 4 shows the effect of two different link failures.

Figure 4: Small Loops Topology After Link Failures



4 Large Loops Topology

This topology is composed of many large loops (more than two) of stations connected to a single backbone switch. Figure 5 illustrates the topology with five stations per loop.

This topology requires STP support from the field switches. The number of hops varies between three and five depending on the devices communicating. The backbone must be assigned the highest bridge priority so that it is the root switch. The loops will automatically 'break' evenly because of the homogeneous link cost.

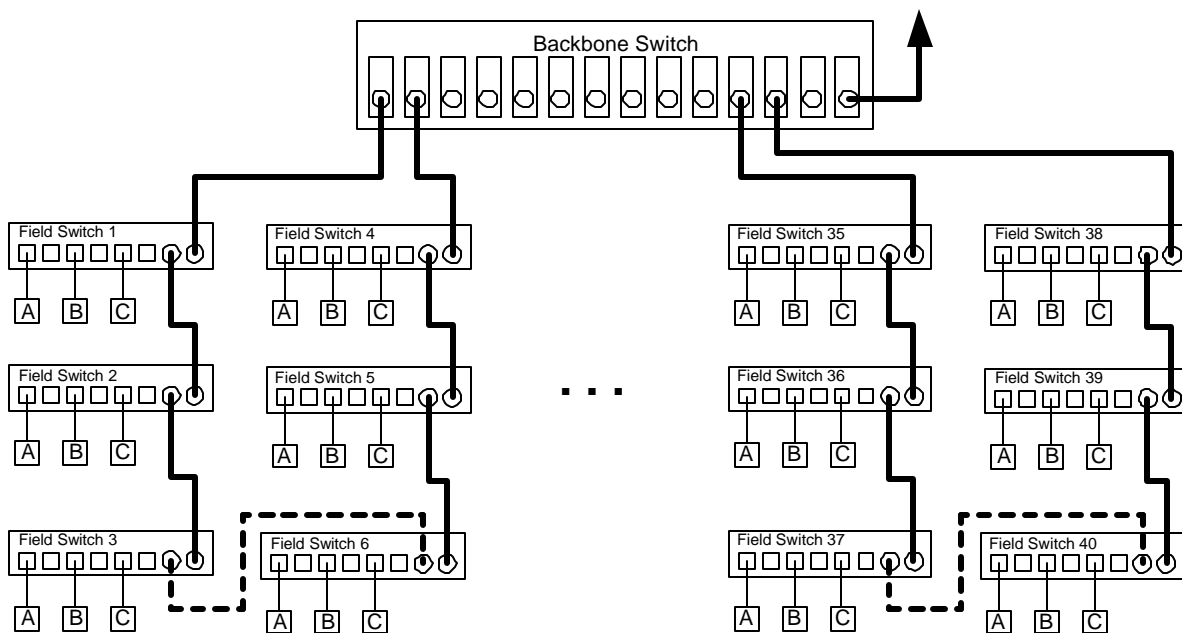
This topology will work with unmanaged field switches however the maximum number of hops between any pair of devices is eleven. This may result in unacceptable latency and allowable station bandwidth.

The maximum sustained bandwidth acceptable from a station is 33 Mbps.

The total cable distance is: $(2/6 * nS * dB) + (5/6 * dS) < 2 * nS * dS$

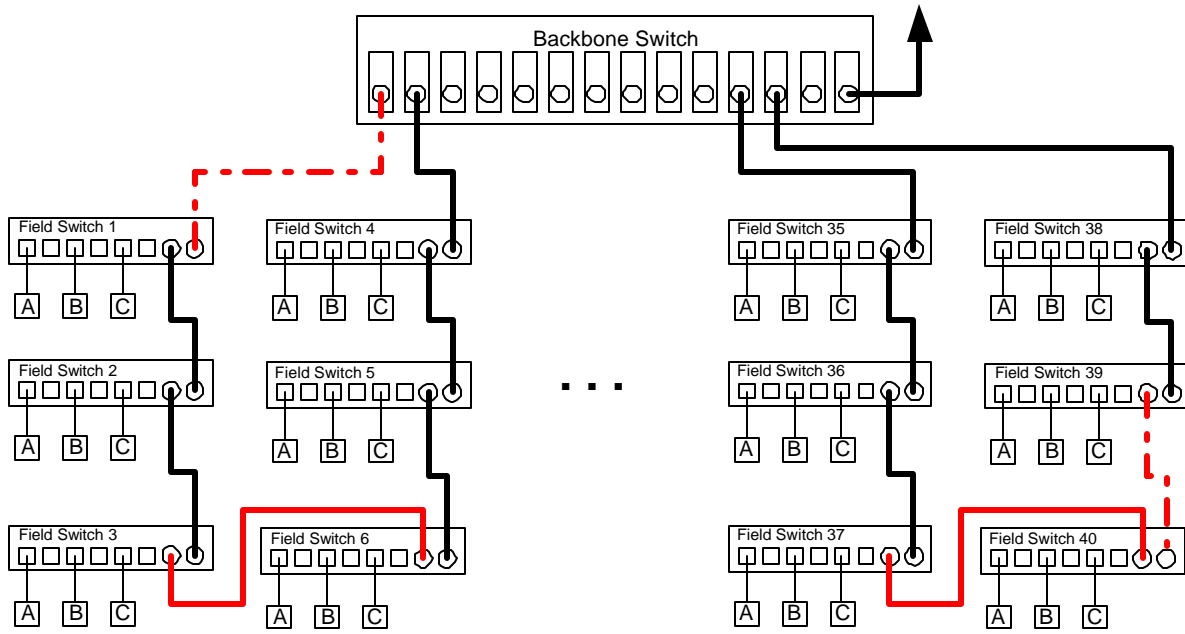
The main benefit of this topology is minimized cable cost. It may provide the simplest cable routing depending on geographical arrangement of the stations. The penalty is increased latency and less available bandwidth per station.

Figure 5: Large Loops Logical Topology



A failure of any single link in this topology is recoverable via STP. Failure of two links on a loop will isolate devices on that loop. Failure of the backbone will isolate all loops from each other. Figure 6 shows the effect of two different link failures. Note that any link failure will result in non-optimal number of hops between end devices.

Figure 6: Large Loops Topology After Link Failures



Before the link failure traffic from A-FS1 to A-FS35 required three hops. After the link failure, that same traffic requires eight hops.

5 Mesh Topology

The mesh topology consists of a single backbone switch; each station has a direct connection to the backbone and a connection to more than one neighboring field switch. Figure 7 illustrates. The inter-station connections result in link redundancy and are routed with more concern about ease of cable routing than the network topology. Some stations may connect more than two other stations if this eases wiring. Figure 9 illustrates a possible physical layout of the inter-station links.

This topology requires STP support from the field switches because of the numerous loops passing through the field switches. The backbone must be assigned the highest bridge priority so that it is the root switch. The inter-station links will then be disabled because every station is directly connected to the backbone. In other words, during normal operation, all traffic goes through the backbone and no traffic goes through the inter-station links.

The mesh topology has many performance benefits including:

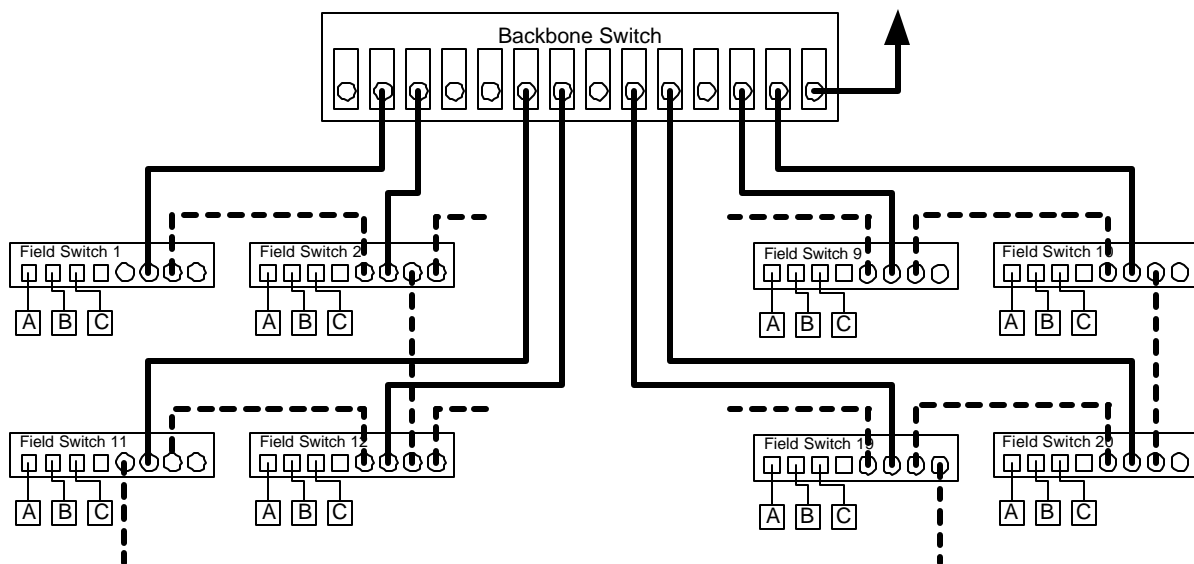
- can tolerate multiple link failures
- can tolerate complete failure of the backbone
- maximum bandwidth available to all stations

The drawbacks are:

- higher cabling costs than the loop topologies
- higher field switch costs since they must be managed and must have four single mode fiber optic ports.
- network complexity

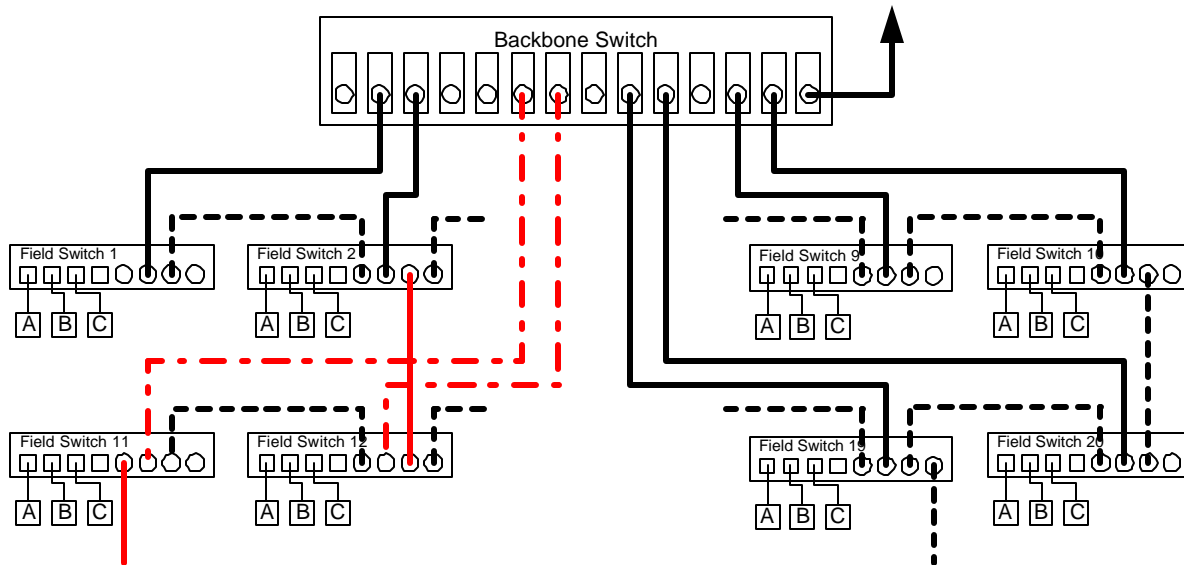
Cabling cost is: $nS * dB + nS * dS = 4 * nS * dS$

Figure 7: Mesh Logical Topology



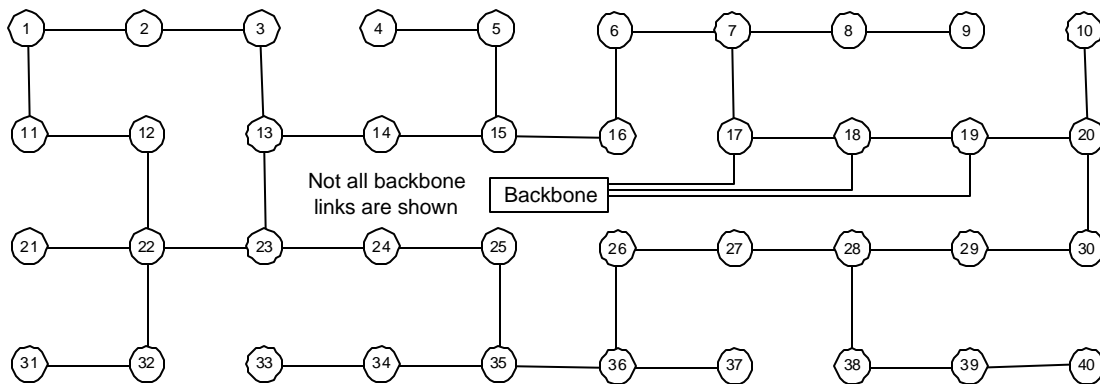
The additional complexity of the mesh topology makes it more difficult to see exactly how the topology changes after a link failure. It depends entirely on how the bridge and port priorities are configured on the field switches. This does require a networking professional to properly configure the network to get the desired fail-over behavior. Figure 8 shows a possible resulting topology after a pair of link failures.

Figure 8: Mesh Topology After Link Failures



Before the link failure, traffic from A-FS12 to A-FS9 required three hops. After the failure, that same traffic requires four hops.

Figure 9: Mesh Topology Physical Layout



6 Dual Backbone Topology

The dual backbone topology has two backbone switches with each field switch having a link to each backbone as shown in figure 10.

The dual backbone topology has the following benefits:

- can tolerate complete failure of a backbone without loss of uplink facilities
- maximum bandwidth available to all stations
- does not require managed field switches

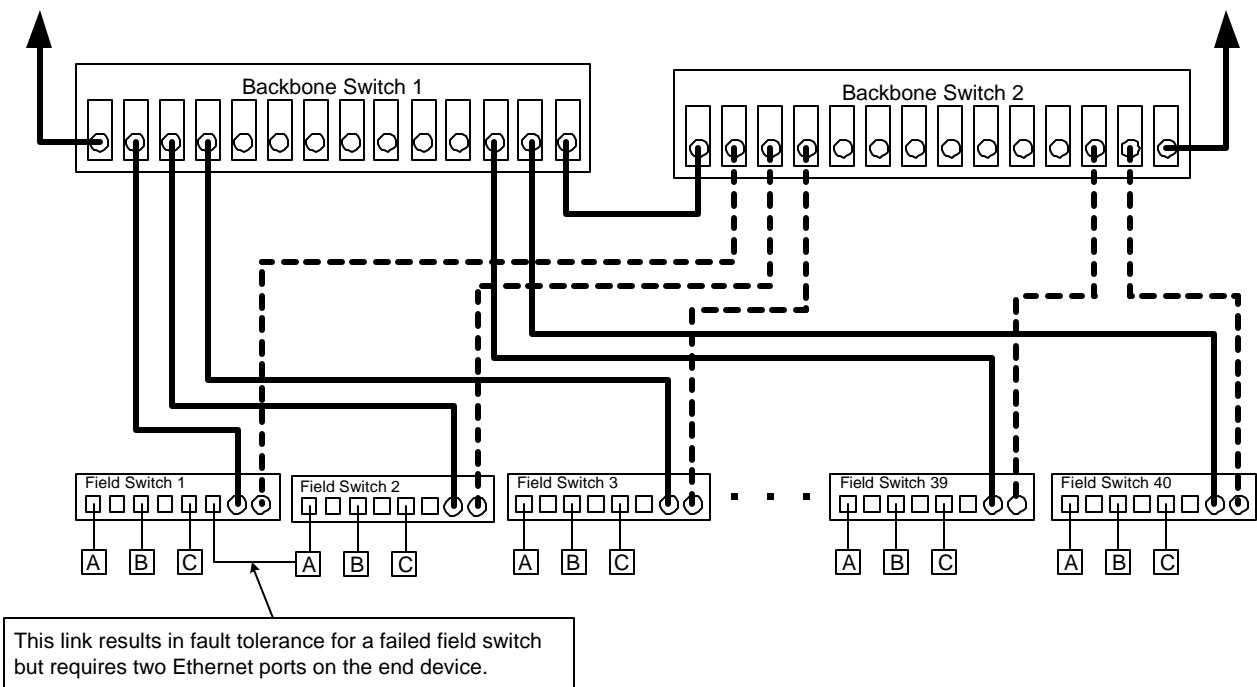
The drawbacks are:

- greatest cabling cost of all topologies
- requires two backbone switches
- cables from field switches to backbones should be routed to reduce likelihood of both cables failing simultaneously

It is worth noting that none of the topologies mentioned so far are tolerant of a field switch failure. To accomplish this requires that each end-device have two Ethernet ports with each port being connected to a different field switch. Coupled with the dual backbone topology yields a highly fault tolerant network.

Cable length is: $2 * nS * dB = 6 * nS * dS$

Figure 10: Dual Backbone Logical Topology



7 Comparison of Topologies

The following table illustrates the relative merits of the topologies discussed in this paper. These figures apply under normal conditions i.e. no link failures. All the topologies will have reduced performance after link failures.

Criteria	Small Loops	Large Loops	Mesh	Dual Backbone
Tolerant of Single Link Failure	Yes	Yes	Yes	Yes
Tolerant of Backbone Failure	No	No	Yes(1)	Yes
Range in Switch 'Hops'	3-5	3-7	3	3
Relative Worst Case Frame Latency	1.66	2.33	1.0	1.0
Max. Bandwidth per Station (Mbps)	50	33	100	100
Relative Cabling Cost(2)	1.75	1.0	2.0	3.0
Requires Managed Field Switches	No	Yes(3)	Yes	No
Number of S.M. Ports per Field Switch	2	2	4	2
Number of S.M. Backbone Ports	40	16	40	80

1. Uplink to management facilities is lost.
2. Estimate for grid model; highly dependent on geographical layout of stations
3. Can be done without managed switches however performance suffers.

8 Calculation of Latency

A real-time control network may be very sensitive to the communication latency between end devices. Following is a calculation for latency based on the number of switch hops. The number of hops is relevant because the majority of switches today are store and forward. Store and forward results in a minimum latency of one frame time per hop since the entire Ethernet frame is stored by the switch before it is forwarded out the destination port.

The latency for a 1500 byte frame is roughly 125 us at 100Mbps. With traffic congestion in the LAN, each field switch may queue up frames creating even more latency. For sake of argument, let the field switch have a nominal queue with eight frames (a very busy LAN!). Latency is now 1ms for the switch. The one-way latency is then the number of hops times one millisecond.

The worst-case round-trip latency for the large loops topology is thus 26ms. This is still much less than the timeout of higher layer protocols like TCP.

It should be noted that the rule of thumb is to allow no more than seven hops between any pair of end devices.

5.0 Conclusions

Fault tolerance comes at a cost; this is no big surprise. As with all engineering tasks one must compromise between performance and cost. The geographic arrangement of the distributed stations plays an important role in the choice of the network topology as does the performance requirements and need for fault tolerance. All these factors must be considered along with the total cost before making a final decision as to the best topology for a given application.

References:

1. The Switch Book, Rich Seifert, Wiley
2. ANSI/IEEE Std 802.1D, 1998 Edition