Understanding and Validating Ethernet Networks for Mission-Critical Protection, Automation, and Control Applications

Saroj Chelluri
NTPC Limited

David Dolezilek, Jason Dearien, and Amandeep Kalra
Schweitzer Engineering Laboratories, Inc.

Original edition released March 2014

Understanding and Validating Ethernet Networks for Mission-Critical Protection, Automation, and Control Applications

Saroj Chelluri, *NTPC Limited*
David Dolezilek, Jason Dearien, and Amandeep Kalra, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—The communications standard IEC 61850-5 identifies fast messages that perform high-speed automation, protection, and interlocking to meet or exceed a transmission time of 3 milliseconds as Type 1A, Performance Class P2/P3. Modern microprocessor-based devices and Ethernet networks routinely meet this requirement when everything is working as expected. One of the most important acceptance criteria (and perhaps least understood) is the maximum transmission time when unexpected things do happen and messages are delayed. Because not all paths in an Ethernet network perform the same, this paper introduces path performance classifications that illustrate the minimum and maximum transfer times between two devices.**

The telecommunications performance standard IEC 60834-1 is commonly used to evaluate point-to-point high-speed automation and interlocking. It describes the overall operating time between the instant of the change of state at the command input on the source device and the instant of the change of state at the command output on the destination device. This includes propagation time and any additional delays. IEC 60834-1 further defines transmission dependability as the ability to receive each command message within the fixed actual transmission time defined by the application, in this case 3 milliseconds.

IEC 61850-5 specifically states that testing and verification of the complete transfer time must be performed during site acceptance testing using the physical devices and network equipment. Methods to test and validate message transmission during normal Ethernet packet delivery as well as during path failure are introduced in this paper based on both Rapid Spanning Tree Protocol (RSTP) and Parallel Redundancy Protocol (PRP).

I. Introduction

Modern microprocessor-based protection, control, and monitoring (PCM) intelligent electronic devices (IEDs) perform many functions and communicate data related to these functions. Communications-assisted PCM automation and control schemes that require high speed and high availability rely on mission-critical communications networks. Robust real-time mission-critical communications networks and digital messaging, in turn, require appropriate engineering design and validation. This paper contains illustrations of test methods in IEDs and test devices to verify the performance of Ethernet networks as previously described in the paper "Design and Validation Practices for Ethernet Networks to Support Automation and Control Applications" [1].

A data channel is the combination of a method of conveyance and a communications medium for the purpose of moving data from their source to the destination. IEDs have evolved to make use of many different types of data channels to communicate protection, control, metering, monitoring, automation, digital fault recorder (DFR), event report, and settings information. For this paper, we focus on mission-critical signaling used to remotely assist or accelerate automatic control functions. The speed and accuracy of the response of a control, automation, or protection system are proportional to the signaling channel speed between IEDs.

Early signaling methods included hard-wired contacts, dedicated pilot wires, and leased telephone lines to convey on and off signals. Signal information was also superimposed on carrier frequencies over power line carrier. Modern signaling is performed via a communications channel composed of digital messaging and a communications medium. The availability of the communications medium and the behavior of the digital message together define the required speed and reliability of the signal channel. Different types of messages are published one at a time into a network through both serial and Ethernet interfaces, providing one or more types of information.

Serial communications channels are usually installed as direct, single-purpose connections following a single path between the IED and the data receiver, such as another IED or computer. Ethernet, in turn, is deployed as a shared media interface within the IED, where signaling, control, and settings channels simultaneously share the same interface. IEDs divide each digital message stream into Ethernet packets, with each packet containing destination information and payload. The Ethernet network includes perimeter Ethernet switch ports cabled to IEDs and backbone ports used to cable Ethernet switches, alternately referred to as Ethernet bridges, to each other. The Ethernet signal channel between two IEDs includes the IED Ethernet interface cabled directly to perimeter ports on Ethernet switches as well as the cables and switches between. The design and validation challenge presented by Ethernet is that every pair of perimeter ports can and will have multiple combinations of backbone cables and switches between them and that every combination has to be understood and verified. Data paths through multiple cables and switches between perimeter ports change as failures occur and new logical paths are created to deliver data through different cable and switch combinations.

TABLE I
IEC 61850 ETHERNET TRAFFIC RECOMMENDATIONS [2]

Function Type	Message	Protocol	Maximum Delay (ms)	Bandwidth	Priority	Application
1A – trip	Generic Object-Oriented Substation Event (GOOSE)	Layer 2 (L2) multicast	3	Low	High	Protection
1B – other	GOOSE	L2 multicast	10 to 100	Low	Medium high	Protection
2 – medium speed	Manufacturing Message Specification (MMS)	Layer 3 (L3) Transmission Control Protocol with Internet Protocol (TCP/IP)	<100	Low	Medium low	Control
3 – low speed	MMS	L3 TCP/IP	<500	Low	Medium low	Control
4 – raw data	Sampled Value (SV)	L2 multicast	4	High	High	Process bus
5 – file transfer	MMS	L3 File Transfer Protocol (FTP) via TCP/IP	>1000	Medium	Low	Management
6 – time synchronization	Time synchronization	L3 IP (Simple Network Time Protocol) L2 (Precision Time Protocol)		Low	Medium high	General phasors, SVs
7 – command	MMS	L3 IP via TCP		Low	Medium low	Control

The signal application is used to accomplish communications-assisted functions, and the Ethernet signal method delivers data as packets between devices. Signal data latency is defined as the time duration for data to travel from the source IED to the receiving IED. The IEC 61850 Communication Networks and Systems for Power Utility Automation – Part 90-4: Network Engineering Guidelines technical report defines latency of communication as the delay between the instant that data are ready for transmission and the moment they have been completely received at their destination(s) [2]. IEC 61850-5 describes the IEC 61850 traffic recommendations [3]. It does not identify the other necessary traffic on the IEC 61850 Ethernet network for maintenance, telephony, video surveillance, and so on.

From the recommendations for IEC 61850 Ethernet traffic in Table I, we see that GOOSE used for protection has the highest priority and the shortest maximum delay. Control blocking schemes, via GOOSE or any other method, require a 99.99 percent success rate and direct control schemes require a 99.999999 percent success rate of receipt of digital messages. Direct tripping, via delivery and processing of a GOOSE or other message, is typically expected to occur within 20 milliseconds [4]. Failure is defined by the absence of the message at the receiving end or, for direct control, a delay in delivery greater than 18 milliseconds. Therefore, IEC 61850 Type 1A, Performance Class P2/P3 requires that the system meet the 3-millisecond transmission time 99.9999 percent of the time and have a delay no longer than 18 milliseconds for the remainder. This means that the new design challenge is to create an Ethernet system that delivers packets quickly and reliably with minimum additional delay due to the failure of an Ethernet interface, cable, or switch. This requires high device reliability to keep the path failures to a minimum. System availability analysis based on IEC 61850-5 measures of reliability predicts the ability of each system to meet IEC 60834-1 dependability and security requirements. Network reconfiguration around a path failure is required to be fast enough to satisfy the maximum packet latency during the failure. And, when this cannot be satisfied with the chosen switch network, redundant networks and redundancy protocols may be required.

Messages can include a large amount of information distributed among multiple packets, and the duration is from the first through the last packet. Signals can also be published in a single packet that is repeated. If the first several packets are not delivered, the signal data latency duration includes the time between the publication of the first packet and the first successful delivery of a message packet. Latency validation requires identification of which cables and switches affect a specific channel so that their behavior can be measured, verified, and improved when necessary. Once the initial channel is tested, additional channel paths that result from rerouting packet delivery around a failure must be tested. Even in a small network with a source IED on one switch and a destination IED on another switch, there will be a very large number of switch and cable failures to simulate and test. As an example of the size of this task, consider a simple Ethernet network made of a ring of ten switches, as shown in Fig. 1.

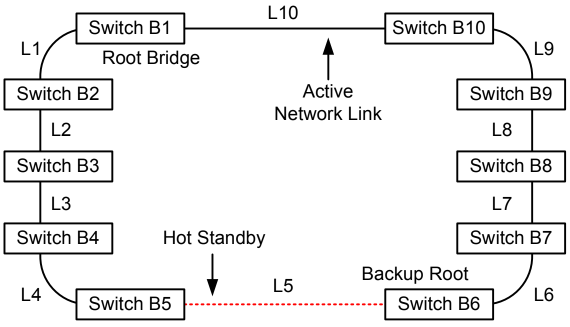

Fig. 1. Ten-Switch Ring Network Failure Combinations

For every possible path among the ten switches, there are two switches directly connected to the IEDs and eight that are not directly connected to the IEDs. There are 45 possible perimeter port source switch/destination switch pairs, as shown in (1). There are 20 possible failure points (ten switches plus ten links), resulting in 900 link and switch failure combinations, as shown in (2) using the combination formula C(n,r).

$$C(n,r) = \frac{10!}{8! \bullet 2!} = 45 \quad (1)$$

$$\begin{aligned}&(C(n,r))(\text{Number of Switches} + \text{Number of Links}) = \\ &45 \bullet 20 = 900\end{aligned} \quad (2)$$

where:

n is the number of possible links between the switch pair being tested.

r is the number of possible switches between the switch pair being tested.

Detailed analysis of which failures actually affect data flow between the IEDs for each path determined that 285 failure scenarios actually interrupt the path between the two IEDs. Testing is required to understand the latency introduced due to reconfiguration of these failure modes. Specific knowledge of how individual switches work during failure modes is required to know which paths to test. If it is not known or clear which cables and switches will affect the channel, all 900 scenarios must be tested. Test results verified that all perimeter ports on the switches have the same performance, so it does not matter which port is used to connect to the IEDs.

II. Ethernet Networks for PCM IEDs

The Ethernet communications system is a switched network with several physical cable paths or loops, like cables in an electrical distribution system, where one is active and the others may act as hot standby. Ethernet packets are prevented from traveling in a loop back toward their IED of origin by an Ethernet switch mechanism that virtually opens and stops the packet flow. This works much the same way electric energy is prevented from looping back toward its source by a power system switch that physically opens the circuit and stops energy flow. And, similar to an automated energy distribution network, when there is a failure in the network, the system detects and isolates it and then reconfigures among the hot-standby paths to quickly begin delivering packets again. When a portion of the Ethernet network is unavailable to deliver packets, we refer to it as being dark. Therefore, the period of time a network channel is interrupted and cannot deliver packets between perimeter ports is referred to as network darkness. Similar to energy distribution systems, it is extremely important for signaling to keep periods of network darkness short and infrequent in Ethernet packet distribution networks.

For energy distribution, the system interruption duration includes both network reconfiguration and reestablishment of energy delivery to the consumer. For Ethernet packet delivery systems, the interruption duration similarly includes network reconfiguration plus the subsequent reestablishment of the channel and packet delivery to the consumer. The IEC 61850 Communication Networks and Systems for Power Utility Automation – Part 90-4: Network Engineering Guidelines technical report echoes common best engineering practices in requiring that the system average interruption duration, or period of darkness, for each possible Ethernet packet delivery channel failure mode be tested and measured [2]. Best engineering practice is to design systems where the median time of darkness is brief enough to be within the maximum signal delay time. The statistical distribution indicates the worst-case signal channel delay caused by the failure of an Ethernet switch. The duration and probability of the worst-case delay must be understood and mitigated via selection of network devices with sufficient availability, measured as mean time between failures (MTBF) in years.

The time duration to perform PCM signaling includes processing within both the source and destination IEDs as well as propagation of the digital message through the network. Overall application reliability is maximized via dual primary PCM applications, each with its own digital messaging network. Testing methods presented in this paper are equally applicable to testing individual or dual primary networks. Even though both serial and Ethernet networks can be deployed individually or redundantly, it is not possible to answer questions about Ethernet network behavior the same way it has been possible with serial networks. For example, multiservice Ethernet shares the available bandwidth with signaling and other protocols, which may affect message delivery behavior. Also, message parameters in the Ethernet packets work in concert with switch settings to control signal channel paths, and therefore delivery performance, through the network. Perhaps the most useful difference Ethernet provides is the ability to reconfigure after a cable or switch failure to use the hot-standby path. Once reconfiguration is completed, signaling proceeds normally; however, periods of darkness during the reconfiguration may impact the signaling during a power system event. These differences, which make Ethernet networks flexible for reconfiguration after failures, create a challenge for understanding Ethernet signal channel behavior.

III. Signal Transmission, Transfer, and Transit Time

The transfer time specified for an application is the time allowed for a signal or data exchange through a communications system. Transfer time is shown in Fig. 2 (which is from IEC 61850-5) as the time duration between the action of communicating a value from the logic processing of one device to the logic processing within a second device as part of an application. The time duration to publish signal information from Physical Device 1 (PD1), deliver it via a protocol message, and act on it in Physical Device 2 (PD2) is the transmission time of the signal or information. This transmission time duration represents actually performing an action as part of a communications-assisted automation or protection scheme. The transit time, t_b, is the time duration for the message to travel through the communications network.

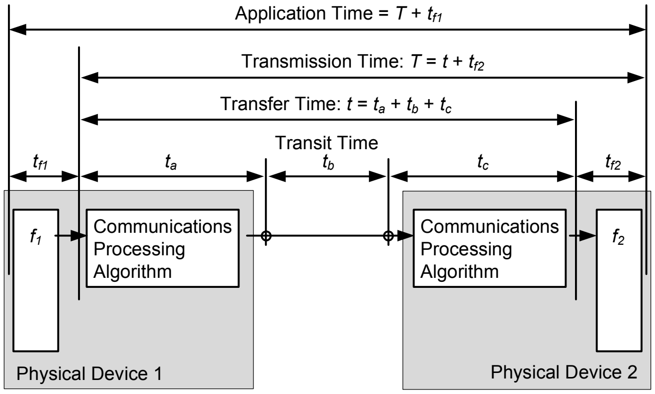
Application Time = $T + t_{f1}$
Transmission Time: $T = t + t_{f2}$
Transfer Time: $t = t_a + t_b + t_c$
Transit Time
t_{f1}
t_a
t_b
t_c
t_{f2}
f_1
Communications Processing Algorithm
Communications Processing Algorithm
f_2
Physical Device 1
Physical Device 2

Fig. 2. Transmission Time and Transfer Time Based on IEC 61850-5

Enhancements to IEC 61850 documented in Edition 2 numerate different types of messages and their associated transfer times, as shown in Table II.

TABLE II
IEC 61850 TRANSFER TIME REQUIREMENTS [2]

Transfer Time Class	Transfer Time (ms)	Application Example
TT0	>1000	Files, events, log contents
TT1	1000	Events, alarms
TT2	500	Operator commands
TT3	100	Slow automatic interactions
TT4	20	Fast automatic interactions
TT5	10	Releases, status changes
TT6	3	Trips, blockings

Questions that are answered via direct and obvious test procedures for serial networks are no longer easy to answer using Ethernet. Some of these questions include the following:

- Question 1: How do I validate the time duration between a power system event and a subsequent mitigation reaction in a remote IED (the total signal application time) via an Ethernet signal application?
- Question 2: How do I validate the transmission time duration between detecting an event in one IED and a subsequent mitigation reaction in a second IED?
- Question 3: How do I validate the transfer time duration between publishing a message in one IED and subsequent message processing in a second IED?
- Question 4: How do I validate the transit time duration of delivering messages between IEDs?
- Question 5: How do I verify the impact of a failure and reconfiguration to a hot-standby Ethernet network path for each of the previous questions?
- Question 6: How do I verify that the Ethernet switches are configured properly for the signal message parameters?
- Question 7: Will the signal channel be affected if I expand the network?
- Question 8: Will the channel reliability increase with the use of Parallel Redundancy Protocol (PRP)?

For each of these questions, network, protection, and automation engineers often ask: How would I know during the design phase? How would I know during a factory acceptance test? How would I know during an acceptance site test? How would I know as part of ongoing monitoring? Many other questions about the IEDs, protocols, and Ethernet message configurations are equally important to signaling (though out of the scope of this paper). Signaling via digital messages requires that specific best engineering best practices be used during specification and design [2] [4] [5] [6]. Best engineering practices for test and validation are within the scope of this paper.

IV. IEC 61850 GOOSE FOR AUTOMATION AND CONTROL SIGNALING

There are many Ethernet messages used for signaling purposes, such as MIRRORED BITS® communications tunneled over Ethernet, EtherCAT, IEC 61850 GOOSE, and network global variable protocols. Our testing focused on the internationally standardized IEC 61850 GOOSE message. IEC GOOSE messages used for signaling are most often deployed among other IED Ethernet protocols on a switched Ethernet network and are multicast to multiple subscribers. Therefore, GOOSE messages are not typically published at a rapid fixed frequency rate because this creates too much traffic on the shared bandwidth Ethernet network. Ethernet interfaces support multiple protocol standards simultaneously, including Telnet, FTP, Hypertext Transfer Protocol (HTTP), and the previously mentioned signal protocols. IEC 61850 combines several protocols over the shared-use network as well, and GOOSE can be used for numerous applications with differing performance. Therefore, all messages, including GOOSE, must be carefully designed to share the available IED resources and navigate the Ethernet network correctly. Though out of the scope of this paper, GOOSE application design also impacts transfer time [2] [4] [5] [6].

The signal channel is defined as an Ethernet GOOSE message published at a variable frequency over a shared Ethernet communications multipath network between two perimeter Ethernet ports cabled to IEDs. As per the IEC 61850 standard, a GOOSE signal message is published immediately after a change of state and then several additional GOOSE messages are published in a quick burst after the change of state. These additional GOOSE messages are referred to as retransmissions because they retransmit the signal data in case some of the initial GOOSE signals are dropped or delayed. At least one GOOSE retransmission must be engineered to occur slightly later after the event than the worst-case duration of network darkness. After GOOSE retransmissions, repetitive GOOSE messages (referred to as a heartbeat) are published less frequently and are used by subscribers to supervise the health of the channel.

IEC 61850 Standard Part 8-1 describes the use the IEEE 802.1Q virtual local-area network identifier (VID), IEEE 802.1p message priority, and IEEE 802.3 media access code (MAC), which is used as a multicast access code, as GOOSE message parameters. These attributes need to be designed with care and accuracy by the designers of the IED messages and communications networks [4]. These attributes work in a coordinated fashion with the Ethernet switch technologies described in this paper to provide mission-critical performance. Best engineering practice requires that IED message, Ethernet switch, and network configuration be performed with precision to correctly segregate Ethernet signal packets [7]. Segregation in the network reduces the likelihood that a GOOSE message delivery will be affected by other traffic or that it will interfere with nonsubscribing IEDs. Best engineering practice recommendations for multicast message exchange include the following: assign each GOOSE message a matching virtual local-area network (VLAN) and MAC address unique from any other GOOSE message, allow no multicast messages on the network without VLAN, disable all unused switch ports, configure each switch port to block delivery of unwanted messages, and assign high priority to GOOSE messages. These recommendations were followed when performing network testing for this paper.

V. IEC 61850 GOOSE and Ethernet Network Test Criteria

The IEC 61850 Communication Networks and Systems for Power Utility Automation – Part 90-4: Network Engineering Guidelines technical report provides advice on network engineering and commissioning [2]. Section 5.3.17 describes testing and recommends the following: “Once the network has been designed, its compliance to the requirements needs to be tested, first as a design verification, then during factory acceptance tests and finally at site acceptance.” This technical report also requires that during operation, an appropriate subset of the tests continue to monitor the network so as to detect and mitigate failures.

However, it is very difficult to cause the worst-case event for t_a, t_b, and t_c from Fig. 2 simultaneously. Therefore, best engineering practice requires that we test and measure the worst case for each time individually and calculate the total worst case as the aggregate ($t_a + t_b + t_c$). Experience shows that Ethernet switches designed for Ethernet GOOSE signaling typically deliver packets in a normally operating network in well under 1 millisecond. For this paper, we used IEDs and Ethernet network switches that together meet a signaling transfer time of Type 1A, Performance Class P2/P3 of less than 3 milliseconds for an Ethernet network using the primary channel path [4]. Also, these IEDs synchronize to an IRIG-B source with microsecond accuracy, so IED clock synchronization error is negligible and ignored. The multiple devices are synchronized to the same time source and therefore share the same absolute time. Time synchronization and accuracy are important because we use IED time stamps to calculate time durations for test results.

IEDs will perform each task, such as detecting a power system change and subsequently performing logic and logging a Sequential Events Recorder (SER) report, at some point during each operating cycle. The IEDs do not monitor inputs or processing logic continuously but rather are designed with specific processing intervals that determine how often they scan their inputs and process their logic. For the IEDs used in tests with a 2-millisecond operating cycle, a binary input and a Boolean logic variable change of state are recognized only once every 2 milliseconds. Signal trigger events, such as digital input contact closure and incoming or outgoing GOOSE signal bit changes and their subsequent SERs, will happen at asynchronous points during the PD1 operating cycle f_1 shown in Fig. 2. Signal reaction events, such as logic equations, digital output contact closure, and incoming or outgoing GOOSE signal bit changes and their subsequent SERs, will happen at asynchronous points uniformly distributed throughout the PD2 operating cycle f_2 shown in Fig. 2. Using the range rule for standard deviation, we approximate the standard deviation for time-stamp error to be one-fourth of the operating cycle time and the mean as one-half. Therefore, we approximate the typical time for f_1 or f_2 reaction processing to be one-half of the operating cycle duration. Also, the time-stamp values of the SER are accurate to 0 + 0.5 operating cycle duration. Ethernet network failures are tested to validate how they impact time t_b shown in Fig. 2.

VI. Network Latency and Delays

Network latency is the amount of time it takes to deliver a packet (message) from the source device port to the destination device port across the Ethernet network. Every device in the active channel between the source and destination adds some latency, and each individual latency must be considered. In a case where the entire channel consists of managed Ethernet switches, simple mathematics can be used to calculate the minimum latency that will be observed when no failures exist and no frame is delayed by active message transmissions. A 100 Mbps link moves 1 byte approximately every 80 nanoseconds, and a 1 Gbps link moves 1 byte approximately every 8 nanoseconds. Therefore, the latency through any one switch is directly related to the size of the packet, as depicted in Fig. 3.

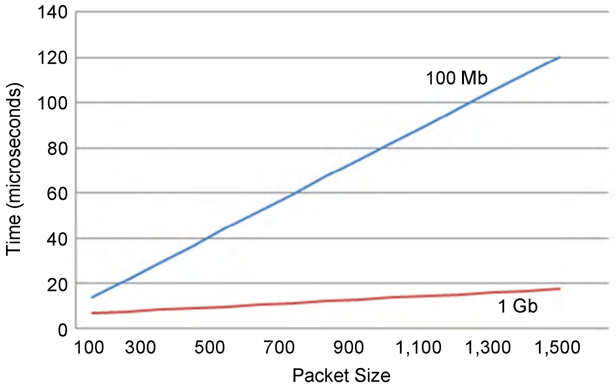

Fig. 3. Ethernet Switch Packet Delivery Behavior Showing Latency Versus Packet Size

Each switch typically adds a delay of 4 microseconds during internal packet processing. Every packet also has an 8-byte interframe gap and a 4-byte frame check sequence, which effectively adds 12 bytes to the size of every frame. Using these numbers, we calculate best-case latency for a packet through a network on a known path. Calculation of the time from the first byte out of the sending IED to the last byte in on the destination IED for a 100-byte packet, plus the 12 additional bytes, that must pass through a possible maximum of 20 switches is as follows:

$$\begin{aligned}&\text{2 links at 100 Mbps} + \text{19 links at 1 Gbps} =\\&\left(2 \bullet \left(112 \bullet 80\right)\right) + \left(19 \bullet \left(112 \bullet 8\right)\right) + \left(20 \bullet 4{,}000\right) =\\&17{,}920 + 17{,}024 + 80{,}000 =\\&114{,}944 \text{ nanoseconds or } 0.11 \text{ milliseconds}\end{aligned} \tag{3}$$

where:

Of the 2 links at 100 Mbps, one is the egress of the sending IED and one is the ingress of the receiving IED.

All switch-to-switch (backbone) links should always use the highest bandwidth available.

The minimum time from when the sending IED starts to put the message on the network to the time the receiving IED receives the final byte and is able to start processing is 0.11 milliseconds. This time is very small in relation to the overall time allotted to the message delivery for the signaling application; however, it does need to be understood.

These calculations do not include any delays that occur while the packet traverses the network, which can be introduced when multiple packets are ready to egress the same switch port at the same time. The packet will be delayed by the time it takes to egress the remaining number of bytes of the preceding packet over the link. If the packet was delayed behind a single maximum-sized packet (1,500 bytes) at every gigabit link in the previously described 20-node example, then the total transmit latency would increase by 230 microseconds (8 • 1,512 • 19).

When two or more packets need to egress the same port at the same time, the network link is considered oversubscribed. Oversubscription is a common and expected phenomenon in packet-based networks and is managed by buffering packets waiting to egress while leading packets are being egressed. Buffering introduces additional packet delivery latency time in relation to the number and the size of leading packets needing to egress the desired port at that given instant. Fig. 4 shows a sample graph of the total cumulative latency at each hop of a packet as it traverses a network, where some hops are oversubscribed and cause extra delay and others do not.

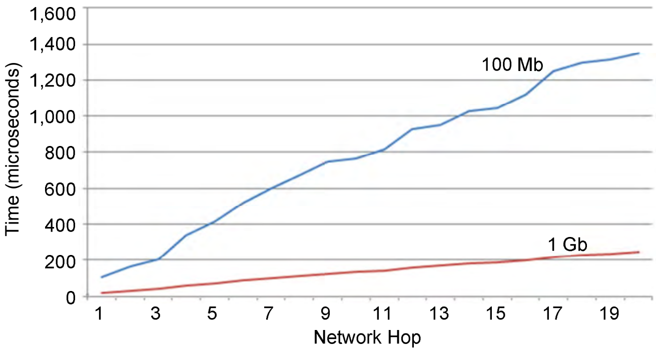

Fig. 4. Ethernet Switch Packet Delivery Behavior Showing Example Cumulative Delay

Buffering memory in the switch is limited, meaning there are limits to the number of packets that can be buffered. If this internal buffering limit is reached, then packets that need to egress a port will be discarded. The discarding of packets due to long-term oversubscription is called saturation. The point at which continuous oversubscription becomes a saturation condition is hardware-dependent, so different devices from different manufacturers will likely behave differently. Fig. 5 shows an example of packet latency on a port. A port that is not constantly oversubscribed will, at times, have longer latency, but when oversubscribed packets egress faster than new packets are buffered, the latency on the port returns to normal. In the case of constant oversubscription, the latency will continue to increase as incoming packets are buffered faster than packets are egressed. When the internal buffers are exhausted, the latency of successful packets becomes constant but packets that can no longer be buffered will be discarded.

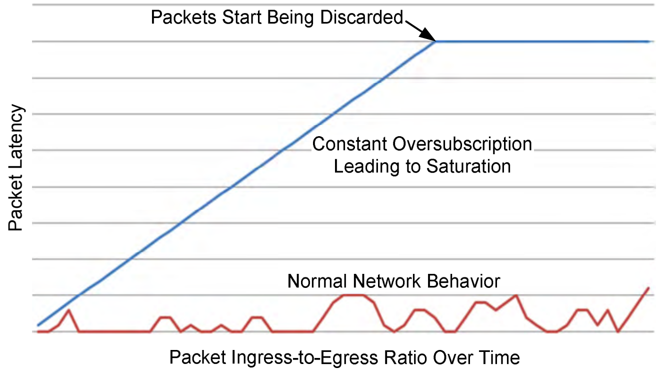

Fig. 5. Ethernet Switch Packet Delivery Behavior Showing Oversubscription and Saturation

When designing a network for mission-critical signal message delivery, it is important to completely understand the traffic flow on the network. This means understanding the type of traffic being ingressed onto the network, the total bandwidth of the traffic, and the frequency that traffic will be put on the network. Understanding these characteristics of the traffic on the network allows for analysis of the possibility of saturation (discarding packets) or possible latency concerns due to large bursts of packets that would not result in saturation but still cause buffering delays. As is shown with the previous calculations, gigabit (or higher bandwidth) links are able to transport a large amount of network traffic very quickly. For example, Fig. 6 shows the bandwidth use of messages being published from three IEDs performing very simple MMS, GOOSE, and minimal other Ethernet-based tasks. Publications from these three IEDs, though normally very low bandwidth, grow in size and frequency, consuming more bandwidth. This quickly saturates a 10 Mbps Ethernet switch port when the IEDs experience a change of state such as a breaker operation. As a result, most switches now support 100 Mbps perimeter ports, and network designers must carefully consider data flow through backbone ports.

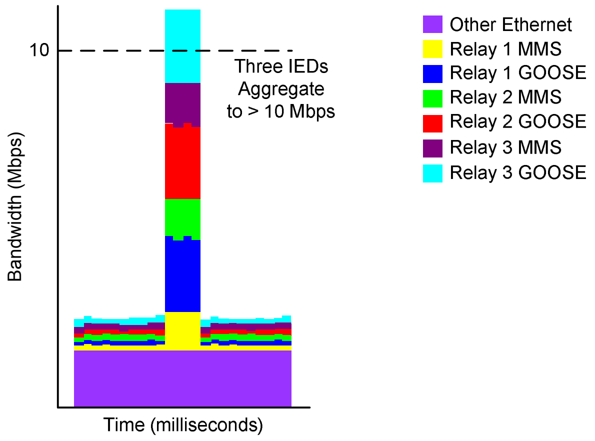

Fig. 6. Bandwidth Use Resulting From Change of State

It is also recommended that critical messages make use of packet prioritization [4] [5]. The IEEE 802.1Q and IEEE 802.1p standards make it possible to apply a priority tag to a packet that enables special handling of high-priority messages by the switches. Depending on the configuration of the switches, it is possible to prioritize the most critical messages so that they will be the next packet to egress the port, jumping in front of all other packets. The only delay to these high-priority packets would be other packets with the same or higher priority.

Special care and consideration must be taken when the bandwidth capacity changes from one network segment to another. For example, having a network with a gigabit backbone will support a large number of devices communicating a large amount of data. If another network segment is connected over a much slower link (T1, for example), then special care must be taken to make sure unneeded traffic does not get onto that slower link because the link will quickly become saturated. Use of the IEEE 802.1Q standard VID allows switches to segregate traffic from these lower-bandwidth segments in order to avoid saturation.

VII. Ethernet Network Reconfiguration

There are several standardized and proprietary algorithms and protocols used to determine primary and failover paths and the rules of how to change between them. There are two types of network failures: switch failures (bridge death) and link failures (link loss). Understanding the behavior of the network reconfiguration algorithm is crucial for engineering a network suitable for critical messaging.

The Rapid Spanning Tree Algorithm (RSTA) is a standard, widely used method that uses the Rapid Spanning Tree Protocol (RSTP) to communicate among switches. When a failure occurs, the RSTA is solved to determine how the network should reconfigure and then RSTP is used to trigger reconfiguration. Parts of the network that are affected by this reconfiguration may be unavailable to deliver packets during the transition or period of network darkness.

The RSTA chooses to always keep the network in the optimal configuration for message delivery, and when a failure occurs, a new optimal configuration is determined and the network transitions to that new configuration. RSTP, by default, chooses active paths (and, in turn, inactive paths) such that the length of all paths among switches between end devices is minimized and uses the highest-bandwidth links possible. It is possible to control these decisions and force specific paths to be active (and others inactive) if required to satisfy engineering needs. If the failure condition is resolved, either by restoring the link that was lost (link restoration) or replacing or fixing the switch that failed (bridge life), the network will revert to the previous configuration that was optimal according to the RSTA. This restorative event will also cause brief network darkness for the same sections of the network that experience darkness during the original failure. It is important to physically wire and properly configure the switches in the network to provide the performance required by the application that will use the network. RSTP allows us to control which switch commands the RSTA of a network by choosing the root bridge using the bridge priority setting. The root bridge of an RSTP-controlled network, often considered the logical center of the network, is very important because all other decisions about active and inactive paths are based on its location. It is recommended that a device with a very high MTBF be chosen for this device and its backup. The backup root is the device that will become the logical center of the network in charge of RSTA decisions in the event the root device fails. A root bridge failure is very traumatic to an RSTP network because all path decisions must be recalculated to use the backup root device.

VIII. Ethernet Signal Application Time Testing

Before testing network performance, it is necessary to verify that the perimeter and backbone ports are configured correctly. For normal operation and every failure mode, each perimeter port must demonstrate correct message egress. This test is performed via a network configuration test device. Every message configuration combination of MAC address and VLAN is injected into the network, and the display shows which messages successfully egress each perimeter port. This answers Question 6 from Section III.

To a small degree, generic surrogate devices and simulation tools have been used to simulate and test Ethernet communications during manually forced link loss and switch failures. However, these results obtained in the laboratory environment require much manual effort and often vary significantly from real-world applications. Also, these tools require a higher level of expertise and understanding of Ethernet-based communications than do the IEDs and switches. More importantly, they are not designed to perform repetitive tests nor are they capable of automatically triggering link loss and switch failures. However, many modern IEDs have built-in communications statistics and logic capabilities that are used to monitor real-time network performance. Using installed IEDs to calculate network latencies and performance parameters provides efficient and constant monitoring of network performance. Also, specialized surrogate devices and extra IEDs installed for test measurements provide easy and accurate measurements.

The total signal application time duration between a power system event and a subsequent mitigation reaction performed by a remote IED (see Question 1) is measured using surrogate synchronized logic IEDs (SLIs). Surrogate devices are attached to laboratory and in-service systems to simulate power system actions and monitor IED reactions for test purposes. These SLIs have high-accuracy synchronization to an IRIG-B time source, have time-synchronized logic, and create high-accuracy digital SER reports. The SLIs trigger logic precisely at the top of the second with microsecond accuracy and, when synchronized to the same time source, will start test activities at precisely the same point in time.

Using synchronized logic, SLI1 in Fig. 7 triggers a simulated power system contingency change of state as a contact output wired to a contact input on PD1 precisely at the top of the second. SLI2 starts a timer at the top of the second. After detecting a contact input, PD1 publishes GOOSE messages with change-of-state data to PD2, which then closes an output contact as a mitigation reaction. SLI2 detects the PD2 output as a contact input and stops the timer as the total signal application time duration. SLI timers are accurate to 0 + 1 millisecond due to a 2-millisecond operating cycle duration and precision starts. For verification, the SLI1 output contact was also temporarily hard-wired to SLI2, and the time duration between the two input contacts on SLI2 was separately measured and confirmed the accuracy of the top-of-the-second timer in SLI2. This means that multiple SLIs can be distributed over any distance and create precise-time measurements via digital messaging alone when synchronized to the same time source. Typical total signal application time was measured to be less than 14 milliseconds.

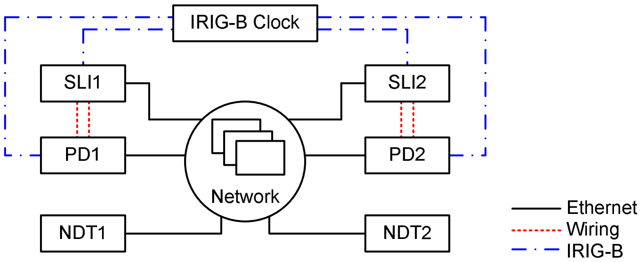

Fig. 7. Test Network

During these tests, the transmission time duration shown in Fig. 2 (which relates to Question 2) is simultaneously measured in both the surrogate SLI and other IEDs. SLI1 records an SER of the change of state as it is published in an outgoing GOOSE message; SLI2 records the reception of that change-of-state bit in a GOOSE message. This transmission time duration measurement is the difference between the SER time stamps and is accurate to 0 + 1 millisecond due to a 2-millisecond operating cycle duration and top-of-the-second precision logic. The second SER calculation of transmission time is the difference between the SER of contact input change of state in PD1 and the GOOSE payload bit change of state in PD2. These IEDs also have a 2-millisecond operating cycle duration. However, these IEDs are protective relays designed to observe power system cycles and are not optimized to start test logic processes at a precise time of day. Therefore, both the timer start and the timer stop in PD1 will each have a typical error of 0 + 1 millisecond, for a total duration calculation typical error of 0 + 2 milliseconds. For each of these tests, the change of state is controlled automatically for repetitive testing of large numbers of samples or by a front-panel pushbutton on an IED for in-service samples. Typical transmission time was measured to be ≤ 4 milliseconds for SLIs and ≤ 5 milliseconds within protective relays.

The transmission and transfer time tests were performed by coordinating the change of state with the network failure to confirm typical times. Then the failures were tested separately in an automated fashion to obtain a statistically significant number of samples in order to understand the statistical distribution, mean, median, and standard deviation.

The transit test requires injecting specific test packets and should not be continuously run on in-service systems. However, the transmission and transfer time tests can be performed in a laboratory, during a factory system test, during a site acceptance test, and continuously as a system self-test function. They are performed in devices executing the control and automation applications and in surrogate devices added to the system specifically for test and validation. Once tested for a specific application, IEDs will perform similarly in the laboratory and in service. However, the times will change with changes in network traffic and path failures. During these tests, the transfer time duration between the SLIs and PDs in Fig. 7 (which relates to Question 3) is simultaneously calculated. Typical transfer time is calculated to be (transmission time − (t_{f2})/2) with an accuracy of 0 + 0.5 operating cycle duration. Additionally, a ping-pong test can be performed, where the second IED reacts by returning a change of state in another GOOSE message and the first IED measures the roundtrip time. Typical transmission time was calculated to be ≤ 2 milliseconds for SLIs and ≤ 3 milliseconds for protective relays. In each test, the network behaves the same way and the duration difference is a result of measurement accuracy differences.

For test criteria, we chose to satisfy the mission-critical signal application of direct trip or control with a typical signal transfer time of less than 3 milliseconds per IEC 61850 Type 1A, Performance Class P2/P3. Based on using IEDs with a 2-millisecond operating time, the typical signal transmission time is less than 5 milliseconds. The maximum signal transmission time of 20 milliseconds was selected to satisfy specific protection and automation schemes [4]. The time delta between maximum and typical transmission times (in this case, 15 milliseconds) becomes the absolute maximum network darkness duration. The true maximum network darkness duration must be short enough to allow a GOOSE transmission to satisfy the maximum signal duration. We configured the IEDs to retransmit at 4, 8, and 16 milliseconds after the initial change-of-state GOOSE message. Therefore, the true and absolute maximum network darkness values are both 15 milliseconds.

If IEDs are not capable of this retransmission pattern, the test may fail. For example, if the retransmission pattern is to transmit at 1, 3, 7, and 14 milliseconds after a change of state, the absolute duration of darkness would overlap all the messages, so the actual maximum would be 13 milliseconds. Alternately, retransmitting at 94, 500, and 1,000 milliseconds after a change of state will never meet the maximum signal transmission time if the initial message is dropped. In this case, if a failure occurs to disrupt delivery of the first message and then network darkness ends within 15 milliseconds, the change of state will not be published to subscribers until 94 milliseconds later.

We also tested a different category of PDs that have a slower operating cycle and do not have count-up timers. Therefore, we used a countdown timer to verify that a GOOSE ping-pong completes within twice the acceptable transfer time. These IEDs also have a 4.17-millisecond operating cycle duration in 60 Hz systems. Therefore, both the timer start and the timer stop in PD1 will each have a typical error of 0 + 2.08 milliseconds, for a total duration calculation typical error of 0 + 4.17 milliseconds. For each of these tests, the change of state is controlled automatically for repetitive testing of large numbers of samples or by a front-panel pushbutton on an IED for in-service samples. Typical roundtrip transmission time was measured to be ≤ 16 milliseconds, averaging to a one-way transmission time of ≤ 8 milliseconds within these IEDs.

These same transmission, transfer, and transit time duration tests are performed while injecting additional traffic into the Ethernet network. If the traffic does not travel on the specific perimeter ports, any timing differences are a result of the performance of the channel. When additional traffic is allowed on the perimeter ports, it also affects processing in the IEDs.

IX. Ethernet Network Reconfiguration Time Testing

Accurate testing of network darkness during a failure or restorative event requires the use of measurement techniques that are analogous to the application of interest. In the case of a critical GOOSE multicast message application, a multicast message test must be used. Using a standard unicast message or a specialized message, such as ping (which is used to test IP network address connectivity), is not appropriate. Signaling network tests must be performed using a multicast message with no IP address, which is the format of the GOOSE message. Using a ping-based tester will not give accurate results for the reconfiguration times of the network for GOOSE message signaling.

Data transit time duration that requires multiple messages (see Question 4 in Section III) is validated with an independent surrogate network darkness test (NDT) device, which publishes messages that mimic the critical application messages at a fixed frequency and monitors their reception. Network darkness that causes dropped packets is observed by counting the number of consecutive undelivered packets. The period of darkness is calculated as the number of packets undelivered due to loss or delay multiplied by the time between publications. For this testing, the NDT device was set to publish a message every 0.25 millisecond.

Darkness measurements indicate the impact of each failure and subsequent reconfiguration to a hot-standby Ethernet network path (see Question 5). These times are then used to calculate the total application impact.

The NDT device automatically controls and measures the network failure event and restoration (both bridge and link failures and restorations) so that a statistically significant number of samples necessary to understand the statistical distribution of each network are measured. These large amounts of accurate data on many different network topologies, the measurement locations on those topologies, and the different failure modes provide necessary network design information. These data about network darkness durations enable analysis for every possible failure scenario of each port pair in the network. With this information, it is possible to find locations in certain topologies that will always satisfy the needs of the application with sufficiently short durations of network darkness during reconfiguration events. It is important to note that some applications consist of numerous signals. Each source and destination port pair must be considered.

X. In-Service Ongoing Testing

Ongoing testing of in-service IEDs is performed both opportunistically when power system events occur and more frequently by adding a test bit to the signal payload. This single bit will not affect the signal performance or protection logic, but it will support the application and network time measurements described previously. By comparing the number and identity of expected messages and received messages, IEDs also calculate the frequency and duration of network darkness. When ongoing application self-testing is preferred, subscribing relays are programmed to provide different indications regarding the health of the network based on the values of calculated times. For example, when it is established that a subscribing IED calculates typical transmission times of 4 milliseconds, an alarm threshold of 6 milliseconds will prevent nuisance alarms based on time excursions due to inaccuracies associated with the 2-millisecond operating cycle. IEDs are programmed to set a flag if transmission times exceed 6 milliseconds, and a second flag is set if times exceed 20 milliseconds. These flags are annunciated in the form of front-panel light-emitting diodes (LEDs) or a message on the front-panel display as well as reported to supervisory control and data acquisition (SCADA) and operator interfaces.

Different test publication patterns are easily triggered by front-panel pushbuttons. For example, one pushbutton will trigger a stream of 100 GOOSE messages consecutively, while another will trigger one GOOSE message every time a pushbutton is pressed. These tests are very helpful in detecting intermittent network problems for in-service IEDs and when the network is approaching its saturation limit. When implemented in the field, these tests provide validation when ongoing self-tests are preferred or required.

XI. Network Considerations When Adding Future Applications

IEEE C37.238 and IEC 61850-9-2 Lite Edition protocols to support Sampled Value process bus messaging represent a large additional use of bandwidth. Sampled Value message publications for protection and metering translate into 5 percent and 12.3 percent of a 100 Mbps Ethernet link, respectively [8]. Unless networks are properly designed with consideration of the future addition of process bus traffic, it will easily saturate the network. IEEE C37.238 Precision Time Protocol has stringent latency requirements to achieve < 1-microsecond accuracy. Because this time synchronization signal is on the same IED interface as that used for messaging, the availability of the time synchronization signal is dependent on the availability of the network.

XII. Ethernet Network Architectures

Even though RSTA and RSTP algorithmically enable and disable links in a topology to remove physical loops in the network and minimize the distance between any two points (balance the network), they must operate within the physical wiring of the network. The physical wiring of the network has a large impact on the performance characteristics in terms of reconfiguration and network congestion. PRP is a data communications network protocol standardized by the International Electrotechnical Commission as IEC 62439-3 Clause 4. It supports connecting each IED to two independent Ethernet networks, which may also support RSTA and RSTP. In this way, while one network is dark, the other will likely not be and the signal transmission will be more reliable. Without a reconfiguration method like RSTA and RSTP, a PRP network will work for only one failure and then become permanently failed. Transmission, transfer, and transit time tests apply to single Ethernet networks and each independent PRP network.

We performed testing and comparisons of ring, dual star, and ladder topologies using RSTP for reconfiguration. These topologies are shown in Fig. 8, Fig. 9, and Fig. 10. These designs use fiber gigabit backbone links instead of copper gigabit ports (copper gigabit ports operate more slowly during reconfiguration). During the testing, it was found that the actual behavior of the dual star topology was not appropriate for signaling. This behavior was previously unknown and only came to light as a direct result of this testing. The two remaining topologies include the ring and ladder. The ladder is so named because the rows of switches look like rungs of a ladder. The ladder performs best, and IEDs are easily dual-connected in failover mode between the two switches on each rung.

As mentioned previously, we selected the network maximum duration of darkness during reconfiguration as 15 milliseconds. Other specific applications need to be tested based on their individual transmission time criteria. Root bridge death is a very troublesome failure because it disrupts the switch commanding the RSTA and causes extended darkness. Root bridge death was measured separately and, as mentioned previously, should be managed via choosing a very reliable switch to keep the probability of failure to a minimum. Questions 5 and 7 in Section III were answered for every topology and every failure scenario, and the results are summarized in Table III.

As mentioned, the results verified unexpected excursions from acceptance criteria in the dual star topology, and therefore, it is not recommended for signaling. This paper presents analysis of Ethernet switches designed and built specifically for GOOSE signaling, and these results do not apply to other switches. Every application has its own failure condition requirements. Thousands of data samples gathered during automated testing revealed that the ladder topology satisfied our criteria of 15-millisecond maximum darkness duration and that the ring topology fell short. If the network darkness requirement was not as restrictive, then the ring topology could be a viable solution as well as other switches less optimized for GOOSE signaling.

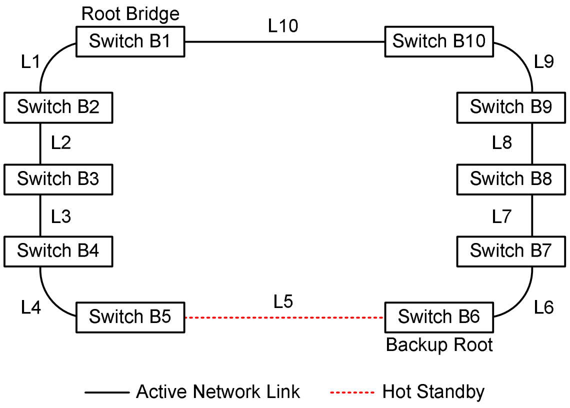

Fig. 8. Ring Ethernet Switch Topology

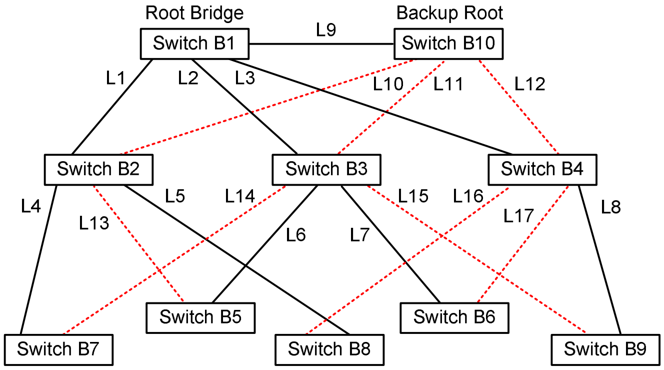

Fig. 9. Dual Star Topology

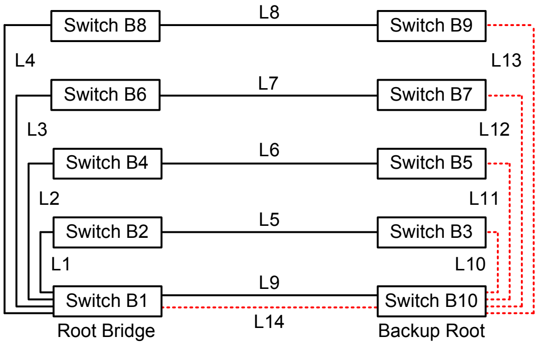

Fig. 10. Ladder Ethernet Topology

TABLE III
RESULTS OF ETHERNET NETWORK RECONFIGURATION TESTS

Topology	Every Channel Meets < 15 ms Maximum Link Loss Recovery Time	Root Bridge Death Typical Reconfiguration Time Is < 15 ms	Non-Root Bridge Death Typical Reconfiguration Time Is < 15 ms	Network Performance Is Unaffected by Additional Switches	Complexity of Choosing Pair of Perimeter Ports That Will Provide Acceptable Signaling Between IEDs
Ladder	Yes	No	Yes	Yes	Port selection does not matter; all pairs are acceptable
Dual star	No	No	No	Yes	Cannot know behavior in advance; we must test each choice
Ring	No	No	No	No	Cannot know behavior in advance; we must test each choice

Testing confirmed that the ladder topology could guarantee acceptable performance (less than a 15-millisecond reconfiguration), regardless of which non-root pair of switches is selected. This guaranteed performance greatly simplifies the task of cabling among IEDs and switches. Values for link failure ranged from 12.8 to 13.8 milliseconds, and non-root bridge loss was always less than 10 milliseconds. Root bridge death was occasionally measured up to 18 milliseconds. This answers Question 8 by showing that when using these switches in a ladder configuration, failover is fast enough to always satisfy the most stringent signaling requirements. Redundancy methods like PRP would not increase the reliability of these switches in a ladder topology but may increase the reliability in other designs.

There are many other benefits to using the ladder topology, including the segregation of network traffic, which reduces latency and saturation concerns. The ladder is simply expanded by adding rungs that will never become part of the original and hot-standby paths of the established channels and will therefore not affect channel performance if they experience failure. This cannot be said for other topologies, such as the ring and dual star. Light travels through fiber at about 186,282/1.467 = 124,188 miles per second. Therefore, latency due to message transit through fiber is negligible for cable lengths within a substation. This means that switches in the field can be configured in the ladder topology regardless of their proximity. Because every non-root switch pair is satisfactory, IEDs can be connected to any perimeter port. This strength and others of both the ring and ladder topologies are listed in Table IV. The dual star topology results were so poor, and characteristics so undesirable, that we chose not to continue considering it for networks performing signaling.

When using a ten-node ring topology, there are some switch pair combinations that have adequate performance (less than a 15-millisecond reconfiguration). A few pairs average well below 15 milliseconds, while other pairs regularly experience network darkness of over 19 milliseconds. However, it is difficult to know which switch pair will always experience less than 15-millisecond darkness duration, so testing is required to confirm channel performance. Once known, appropriate channels are relegated to certain switch combinations in relation to the root bridge. Therefore, this requires that IEDs be connected to specific switches, regardless of their actual physical proximity in the field. Also, as the ring size increases, the network reconfiguration times continue to increase. This means that even though the system may presently meet the critical application messaging needs, it may violate the application timing requirements when expanded. It will be impossible to know in advance when some ring changes will affect performance, and only retesting will verify results. This weakness and others of both the ring and ladder topologies are listed in Table V.

TABLE IV
COMPARISONS OF STRENGTHS OF DIFFERENT ETHERNET NETWORK TOPOLOGIES

Ring Topology	Ladder Topology
Is simple to build. Requires shorter cable runs, which are less expensive. Has maximum IED-to-switch ratio. Only requires two backbone links per switch.	Is very robust and can handle many failures. Has consistent latency in failure conditions. Has consistently small latency. Has very localized network darkness during failure. Can scale without affecting performance. Has localized traffic on network segments. Requires minimum settings changes even for a large network. Has very consistent reconfiguration times. Provides guaranteed locations on network with good reconfiguration times.

TABLE V
COMPARISONS OF WEAKNESSES OF DIFFERENT ETHERNET NETWORK TOPOLOGIES

Ring Topology	Ladder Topology
Has saturation and latency concerns caused by traffic flowing around ring. May require settings changes to every switch in large networks. May have limited maximum ring size. Has variable reconfiguration times depending on the source, destination, and failure location. Only protects against a single failure. Causes failures to impose network darkness onto a larger segment of the network.	May not be as easy to build as a ring. Requires slightly more cabling than a ring (three more cables in the ten-switch topology). Has a slightly smaller IED-to-switch ratio. Requires many backbone speed links on root and backup root switch.

XIII. Conclusion

Simple tools, application and test IEDs, and very specific network test devices play an important role in Ethernet network performance testing. IED features should be deployed for acceptance testing and ongoing monitoring of application behavior, as mentioned in [2]. However, Ethernet network reconfiguration testing requires new special-purpose test devices to verify configuration and performance. These devices must be configurable to use enough resolution and accuracy to measure true performance and automatically trigger link loss and bridge failure to collect statistically meaningful results. Also, they must use appropriate technology to verify network behavior for the specific signal message types, such as multicast GOOSE messages.

Application tests confirmed typical times for an error-free network to be 14-millisecond application, 4-millisecond transmission, and 2-millisecond transfer times. SLIs time-stamp changes and measure these times with an accuracy of + 0 to 0.5 operating cycle duration time. Protective relays time-stamp changes with an accuracy of + 0 to 0.5 operating cycle duration and measure transmission duration with an accuracy of + 0 to 1 operating cycle duration time. These times meet IEC 61850 Type 1A, Performance Class P2/P3.

Reconfiguration tests confirmed that the chosen Ethernet switches, designed specifically for PCM applications, routinely deliver packets with a transit time typically well under 1 millisecond. Network reconfiguration behavior and worst-case transit time depend greatly on the network topology, switch settings, and the design of the switches. Any one of these characteristics can easily mean the difference between meeting the application requirements for critical messaging and failing to do so.

XIV. References

[1] S. Chelluri, D. Dolezilek, J. Dearien, and A. Kalra, "Design and Validation Practices for Ethernet Networks to Support Automation and Control Applications," proceedings of the Power and Energy Automation Conference, Spokane, WA, March 2014.

[2] IEC/TR 61850-90-4, Communication Networks and Systems for Power Utility Automation – Part 90-4: Network Engineering Guidelines, Technical Report, August 2013. Available: http://webstore.iec.ch/.

[3] IEC 61850-5, Communication Networks and Systems for Power Utility Automation – Part 5: Communication Requirements for Functions and Device Models, January 2013. Available: http://webstore.iec.ch/.

[4] D. Bekker, T. Tibbals, and D. Dolezilek, "Defining and Designing Communications Determinism for Substation Applications," proceedings of the 40th Annual Western Protective Relay Conference, Spokane, WA, October 2013.

[5] D. Dolezilek, N. Fischer, and R. Schloss, "Improvements in Synchronous Wide-Area Data Acquisition Design and Deployment for Telecontrol and Teleprotection," proceedings of the 14th Annual Western Power Delivery Automation Conference, Spokane, WA, March 2012.

[6] IEC 61850-8-1, Communication Networks and Systems for Power Utility Automation – Part 8-1: Specific Communication Service Mapping (SCSM) – Mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3, June 2011. Available: http://webstore.iec.ch/.

[7] D. Dolezilek, "Using Information From Relays to Improve the Power System – Revisited," proceedings of the 1st Annual Protection, Automation and Control World Conference, Dublin, Ireland, June 2010.

[8] V. Skendzic, I. Ender, and G. Zweigle, "IEC 61850-9-2 Process Bus and Its Impact on Power System Protection and Control Reliability," proceedings of the 9th Annual Western Power Delivery Automation Conference, Spokane, WA, April 2007.

XV. Biographies

Saroj Chelluri received her BS in electrical engineering and an MBA. She is presently working as general manager in the project engineering division of NTPC Limited. Her job involves the design of auxiliary power supply systems in power plants, including concept designs, preparation of technical specifications, tender engineering, detail engineering, testing, and execution. She has about 25 years of experience in auxiliary power supply system design and execution. She has been extensively involved in medium- and low-voltage system automation designs for the last five years.

David Dolezilek received his BSEE from Montana State University and is a research and development technology director at Schweitzer Engineering Laboratories, Inc. He has experience in electric power protection, integration, automation, communication, control, SCADA, and EMS. He has authored numerous technical papers and continues to research innovative technology affecting the industry. David is a patented inventor and participates in numerous working groups and technical committees. He is a member of the IEEE, the IEEE Reliability Society, CIGRE working groups, and two International Electrotechnical Commission (IEC) technical committees tasked with the global standardization and security of communications networks and systems in substations.

Jason Dearien received his BS from the University of Idaho in 1993. After graduation, he was a founding member of a small startup software contracting business. Later, he was involved in ASIC development at a fabless semiconductor company, working on compression and error correction technologies. In his 12 years at Schweitzer Engineering Laboratories, Inc., he has worked in various product development groups and is presently a senior software engineer in the communications department, focusing on local-area network and security products.

Amandeep Kalra is an automation engineer with Schweitzer Engineering Laboratories, Inc. (SEL) in Pullman, Washington. He has over three years of experience in developing and marketing system-level solutions for the water and wastewater industry, including pump station automation, canal automation schemes, and communications systems. Amandeep worked as a consultant in various technical roles for irrigation districts throughout California and obtained his EIT certification before joining SEL. He has a bachelor of technology degree in instrumentation and control engineering from the National Institute of Technology, India, and a master's degree in electrical engineering from California State University, Northridge.

20140327 • TP6656-01