Making Peace With Communications Networks: What Power Engineers Need to Know About Modern and Future Network Communication for Plants and Substations

Nicholas Seeley
Schweitzer Engineering Laboratories, Inc.

Kurt Concienne
Chevron Energy Technology Company

This paper was presented at the 58th Annual Petroleum and Chemical Industry Conference, Toronto, Canada, September 19–21, 2011, and can be accessed at: http://dx.doi.org/10.1109/PCICon.2011.6085895.

MAKING PEACE WITH COMMUNICATIONS NETWORKS: WHAT POWER ENGINEERS NEED TO KNOW ABOUT MODERN AND FUTURE NETWORK COMMUNICATION FOR PLANTS AND SUBSTATIONS

Nicholas Seeley
Member, IEEE
Schweitzer Engineering Laboratories, Inc.
2350 NE Hopkins Court
Pullman, WA 99163, USA
Nicholas_Seeley@selinc.com

Kurt Concienne
Member, IEEE
Chevron Energy Technology Company
1400 Smith Street
Houston, TX 77002, USA
KurtConcienne@chevron.com

Abstract—Few power engineers seem to be interested in becoming network engineers, yet some of the most recent significant advances in the electric power industry have come in the form of communications-assisted protection and control schemes. Communication is playing a larger role in the electric power conversation, yet there remains a mystique surrounding the low-level details pertaining to how these communications networks operate, which often results in confusion and misinformation being taken as fact. The operational and performance data disseminated to end users are often ambiguous, seemingly focusing on arbitrary figures and numbers that lack proper context and background needed for an engineer to make informed decisions regarding the use of such technology. This paper uses a greenfield offshore platform project as the impetus for investigating the technical details of modern communications networks and their potential impact on power system operations, while navigating some largely misunderstood concepts. The paper discusses the technical foundation of network communication, from the network switch to the intelligent electronic device. This paper addresses issues of latency, redundancy, and failover and how these factors play into modern communications protocols and engineering access traffic. This paper is meant to be used as a guide and reference for engineers looking to understand the concepts surrounding modern communication, while addressing the challenges and complexities common to the use of modern communications networks in the electric power industry.

Index Terms—Ethernet, redundancy, failover.

I. INTRODUCTION

The inspiration for this paper came from a quick test performed on a small Ethernet-networked system in order to validate the failover performance of a network switch that was being considered for use in a power management and load-shedding system for an offshore platform. The particular network switch advertised a 5-millisecond failover time in a simple ring configuration, as shown in Fig. 1. The intelligent electronic devices (IEDs) were configured to send Generic Object-Oriented Substation Event (GOOSE) messages between the two devices, and the receipt of the messages was time-stamped on each end.

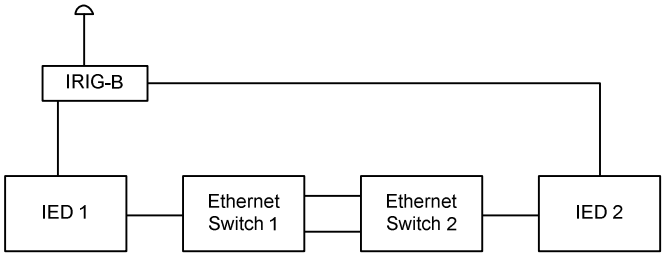

Fig. 1 Test Network

The prevailing wisdom, provided in the switch manufacturer specifications, was that if one path between the two systems was broken, the switch would failover the communication in 5 milliseconds to the working communications link, and this would be observable by looking at the time-stamped sequence of events report inside both IEDs. What was observed was quite different. When the communications link was broken, testers noticed a delay of up to 10 seconds for communication to be reestablished. As this was a hastily thrown-together test setup, the testers recognized that the disparity of the resulting performance compared with the advertised performance could have been a direct result of the faulty contribution of any of the components involved in running the test. The following three culprits immediately came to mind:

- The IEDs may not be adapting properly to errors on the network.
- The switch may not perform to the specifications as stated, or the specifications may mean something other than what the testers expected.
- The testers themselves set up and conducted the test poorly.

The testers decided to research the subject in greater detail to understand what was actually happening at the communications level and find where the observed and expected results diverged. Because the authors are not interested in becoming network engineers and assume that

the audience is likewise uninterested, this paper makes references to texts containing deeper insights into the low-level details but still covers the basics for understanding. It provides context for power engineers to understand the fundamentals of communication when specifying power system protection and control communication at their plants or substations. In addition, the authors conclude with insights they have attained from working on several projects that required the implementation of networked-based communication, specifically focusing on interfacing with groups unfamiliar with power system protection and control communications requirements.

II. NEED FOR RELIABLE COMMUNICATION ON THE PLATFORM

A full power management and load-shedding (PMLS) system is designed to be integrated into the offshore platform, including a complete supervisory control and data acquisition (SCADA) system for engineering access, visualization, remote operation of the power circuit breakers, and automatic generator synchronization, as well as a load-shedding system for maintaining system frequency under loss-of-generation scenarios. Different control schemes within the PMLS system call for various speeds of data transmission and throughput. Engineering access applications often require higher data throughput, but higher-latency protocols are usually acceptable, given the nonreal-time nature of the task being performed. Other data, such as information associated with high-speed load shedding, require ultra-low latency but relatively little bandwidth.

A real-time system, such as a load-shedding system, is dependent upon low-latency communication, and if a portion of the network fails, it is important that as much of the load-shedding system remain enabled as possible. The dependability of communication becomes crucial, and the robustness of the communications system plays an important role, considering that the mean time between failures (MTBF) of an average Ethernet switch hovers around 46 years (as advertised by major Ethernet switch manufacturers). This makes the Ethernet switch one of the weaker links regarding the protection and control aspect of a power system.

Fig. 2 shows a possible configuration of the offshore platform PMLS communications network. The network is configured as a ring, similar to, albeit obviously larger than, the simple test described in Fig. 1. With the data traveling around the ring network, the main question becomes: what is the maximum network downtime permissible? For an essential system, such as load shedding or zone interlocking, acceptable downtime is considerably less than for a simple SCADA system. As such, this platform PMLS system requires minimal downtime, considering the load-shedding function it performs; therefore, the network must be designed accordingly. If the network failover ultimately ends up being in the order of 10 seconds, as originally determined from the initial quick test, the single ring configuration may not be sufficient because the end user may require significantly lower failover times, thereby forcing the use of a different network architecture.

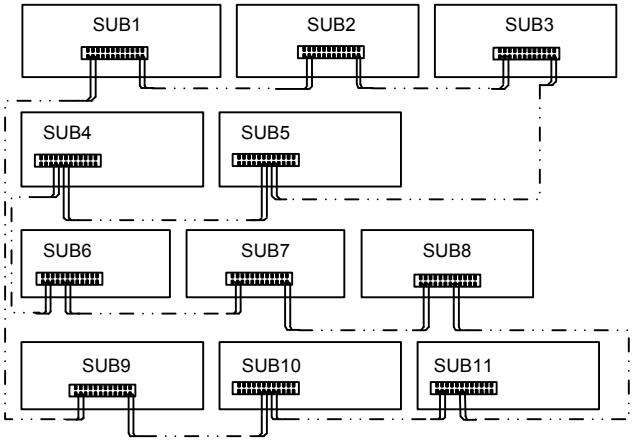

Fig. 2 Simplified Substation Network Connection

The communications network used in modern power systems is becoming a critical part of the overall project. In some cases, such as with zone-interlocking schemes and transfer tripping, the coordination of the protection schemes is dependent on the communications network. Given the critical nature of the communications network for secure and reliable operation of the power system, it seems obvious to want to understand exactly how this modern network communication works, where its weaknesses are, and how modern protocols fit into the picture.

III. ETHERNET COMMUNICATION FOR POWER ENGINEERS

Ethernet communication has erupted on the scene in power engineering and is now becoming more common in a variety of power system areas, including protection schemes, thanks to IEC 61850 GOOSE messaging, time synchronization, Network Time Protocol (NTP), Simple Network Time Protocol (SNTP), and soon to be IEEE 1588. SCADA systems have been using Ethernet communication for many years now, and given the relatively lax timing and determinism requirements concerning SCADA communication (1- to 3-second updates), Ethernet has handled the job reasonably well. However, protection and time synchronization require a level of determinism that SCADA systems do not. Exploring how Ethernet works is key to understanding when and where to apply Ethernet for protection communication.

A. Ethernet and the Physical Layer of the Open Systems Interconnection (OSI) Stack

Most engineers understand the various options that are available when specifying Ethernet communication: 10/100/1000 Mbps bandwidth over copper or fiber-optic cable. The different options highlight the various implementation possibilities of Ethernet, which correspond to the OSI concept of a physical layer. Ethernet, by nature, is an OSI Layer 1, or physical layer, protocol.

In the simplest terms, Ethernet is a protocol. As a protocol, it defines the signal encoding and transmission rate of electrical signals that represent the raw bits transmitted over a copper Ethernet cable or optical light pulses over a fiber-optic

cable. Technically defined as an OSI Layer 1 protocol, Ethernet is described as the physical layer of the network stack and defines how the raw bits are transmitted over their respective media. 10BASE-T Ethernet is different from 100BASE-T, which is different than 100BASE-FX. These are all defined as Ethernet but have different transmission properties. Each is designed to operate over a twisted pair or optical cable, but the encoding of the signal on the line for each is different and, therefore, requires a different physical layer implementation.

B. Ethernet and the Data Link Layer of the OSI Stack

If the physical layer of the communications stack defines the method of communication, the data link layer defines how two devices establish a link in order to exchange data. What is relevant to power engineers, especially when considering IEC 61850, is an aspect of the data link layer called the media access control (MAC) sublayer. The MAC sublayer has an addressing mechanism whereby every Ethernet-enabled device possesses a unique six-hexadecimal octet identifier. Every Ethernet-enabled computer, router, cell phone, and personal data assistant contains a unique MAC address, and this address is used when transmitting data between specific devices. For example, when Device A wants to transmit data to Device B, Device A must address Device B specifically, or it can broadcast data to all devices, and Device B will receive the data.

C. Ethernet and the Network and Transport Layers of the OSI Stack

Internet Protocol (IP) is what most people think of when they think of Ethernet, and it comprises the third layer of the network stack. Most people know what IP addresses are and basically how they are used within the network. Considering that the MAC address is the physical address of the device, the IP address is the network address of the device. Because devices can be moved and placed on different networks (i.e., engineers traveling from site to site and linking to multiple wireless networks), these devices can change IP addresses as needed to establish network communication, but the MAC address of the device remains the same. Amongst other functions, the network layer plays a role in the routing of the packets between networks.

The IP layer is closely related to the transport layer, the final layer of which engineers need to be aware. The transport layer is responsible for the transport and error control of the actual data that are being transferred by the user. Ethernet commonly employs Transmission Control Protocol (TCP) as its transport layer. If the transport layer detects missing, corrupted, or out-of-sequence data, it is responsible for requesting the retransmittal of the data or resequencing the packets. The information contained in the lower layers is necessary to transfer the data to the correct device but is considered overhead information because it has nothing to do with the actual data requested to transfer. Fig. 3 shows a simplified view of the construction of an Ethernet packet.

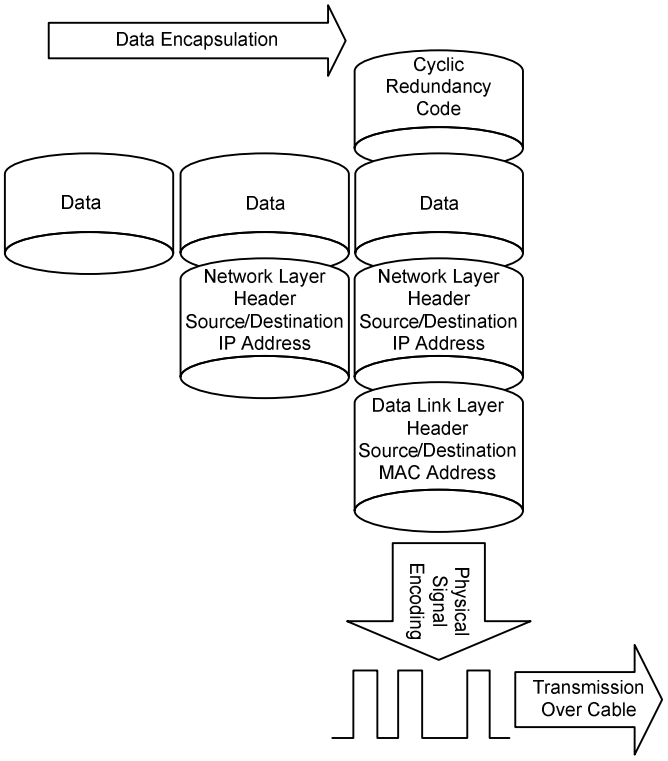

Fig. 3 Simplified Data Encapsulation and Transmission

D. Where Ethernet Switches Fit

As most engineers are aware, the switch provides the physical link between multiple devices. The switch works by reading Layer 2 information from the packet and deciding to which port the packet should be sent. Especially since the advent of IEC 61850, most substation networks employ managed Ethernet switches, which allow certain security functions (MAC address filtering) and traffic control capabilities (prioritization and segregation of data traveling through the network). The switch is able to keep track of which physical devices are connected to which ports on the switch. The switch keeps a MAC address table internally to decide to which port the packet should be sent when a packet is received. It is important to note that the traditional Ethernet switch has absolutely no interaction with data contained in Layer 3 or higher, including the IP address.

E. Putting It All Together

After this brief discussion of the salient points of network communication, an example will help explain the interworking of the layers used at the device level. Staying within the context of presenting information pertinent to power system engineers, there are several behind-the-scenes pieces of the network transaction that remain anonymous for the sake of brevity, and the reader should be aware that the following example provides a high-level overview of the transaction.

Consider the simple network described in Fig. 1. If IED 1 wants to send information to IED 2, the device encodes the data according to the protocol definitions and transmits the data per the physical layer requirements to the switch. This is illustrated in Fig. 3.

Per Fig. 3, the data encapsulation occurs within the IED, as does the encoding of the physical signal. The physical signal is sent across the network cable to the switch. The switch reads the destination MAC address and compares that address with an internal lookup table to determine to which port to direct the packet. After traversing Switch 1 and Switch 2, the packet is received by IED 2, and the data are parsed through the different layers of the stack.

The authors acknowledge that several intermediate steps necessary for this process have been skipped in the explanation and, for the sake of simplicity, ask the reader to accept that certain functions are taken care of internally and are transparent to the user. However, the details provided are a sufficient background for the average engineer working in power system protection.

IV. NETWORK REDUNDANCY

Fig. 1 shows a simple redundant network. Two cables connect Switch 1 to Switch 2. If either the port or the cable fails on one switch, an alternate path to send the data exists. Ring architectures such as this tend to be the most popular for inherent redundancy properties, as well as reduced cabling cost. Note that ring architectures should only be used with managed switches because managed switches have specific properties that allow them to detect a ring and utilize the ring to its fullest potential. Fig. 4 shows a simple ring network.

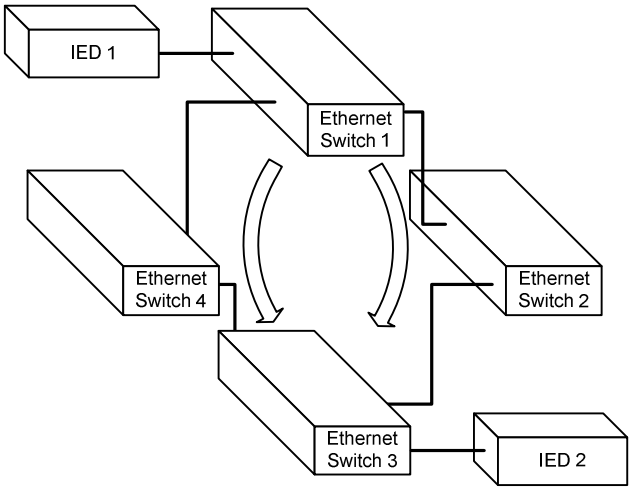

Fig. 4 Simple Ring Network

If IED 1 needs to communicate to IED 2, Switch 1 transmits the data out of both switch ports upon first transmission. When IED 2 responds, Switch 1 makes a decision regarding which path to use for subsequent communication and effectively turns off the alternate port. Communication around the ring only goes in a single direction, so the switch opens the ring, as shown in Fig. 5.

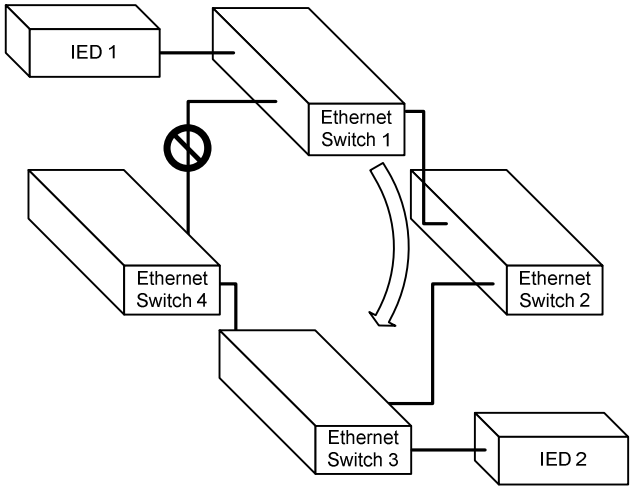

Fig. 5 Communications Path in a Simple Ring Network

When communication is broken by either a switch or cable failing, the switches are responsible for determining where the failure occurred and how to reconfigure themselves to restore service to as many nodes as possible. The time required to do this is often based on the number of switches in the ring, and as stated in the introduction, many manufacturers claim the failover time is in the order of 5 milliseconds per switch. Failover is a function of the specific protocols each switch uses to communicate to each other. Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP) are industry standard methods of network reconfiguration through switches. However, most switch manufacturers offer a proprietary protocol that claims improved performance over STP and RSTP.

V. RING NETWORK FAILOVER PERFORMANCE

The most important issue to recognize is that IEDs operating together over a network operate as a system, and it is the performance of that system that should be scrutinized, because reliable end-to-end communication is the primary goal. When performing such tests, it becomes difficult to unambiguously determine the root cause of a network delay.

A simple test was run with managed network switches from two different vendors, passing information between two IEDs from the same vendor in a configuration identical to that shown in Fig. 1. Vendors A and B were standard ruggedized devices that are typically used in substation environments. Vendors A and B have proprietary protocols that allowed for high-speed reconfiguration of a ring network, both advertised at 5 milliseconds per switch.

This test was done in order to emphasize the notion that the failover numbers specified by the network switch do not

necessarily predict how long it will take the end devices to reestablish communication. An average roundtrip time for IED 1 to send a message to IED 2 and then receive a message from IED 2 in return is roughly 8.5 milliseconds. If each switch takes 5 milliseconds to failover, a total failover time would be somewhere in the neighborhood of 19 milliseconds (Switch 1 at 5 milliseconds + Switch 2 at 5 milliseconds + roundtrip GOOSE protocol latency of 8.5 milliseconds = 18.5 milliseconds).

Table I shows that the performance specified by the network switch may not be reliable enough to give the end user an accurate depiction of complete, end-to-end performance.

TABLE I
NETWORK FAILOVER PERFORMANCE

	Vendor A	Vendor B
Maximum	25 milliseconds	171 milliseconds
Average	21 milliseconds	49 milliseconds

The results recorded in Table I show the average and maximum failover latencies for the two different vendors. It should be noted that the switch for Vendor B exhibited an unexpected method of operation. The failover occurred in less than 50 milliseconds on average, yet when the full ring was reestablished (i.e., the cable was plugged back into the switch to complete the ring), the switch consistently halted communication for approximately 20 to 40 seconds before it resumed normal operation. Essentially, the switch failover was quick, but a return to normal was delayed significantly when the previously broken link was reestablished. This type of behavior could have serious consequences for networks subject to intermittent communications links because of loose or broken connectors.

The authors also tested a standard information technology-grade network switch. The switch did not include a proprietary failover algorithm and relied on a standardized RSTP for failover detection. The RSTP failover algorithm is substantially slower (in the range of 30 to 45 seconds for a failover operation) and does not represent a fair comparison to the other network switches tested, so the results were not included in Table I.

VI. PROTECTION SCHEMES OVER ETHERNET

A. Basics of IEC 61850 GOOSE

After exploring how network schemes are most commonly implemented and how the data travel across a network, it is important to look at the efficacy of Ethernet as a communications transport for power system protection signals. IEC 61850 GOOSE messaging is the industry standard for protection-speed communication over Ethernet, and as such, it is important to understand how IEC 61850 GOOSE makes use of a network.

While this paper does not go into the details of how GOOSE messaging works, a few points regarding the interworking of the protocol are discussed. As many are already aware, GOOSE is a multicast message whose content comprises data predefined by the user and whose broadcast mechanism is based on a change of state of the data within the message or a periodic broadcast time if the data within the data set have not changed recently. Consider the example GOOSE data set in Fig. 6.

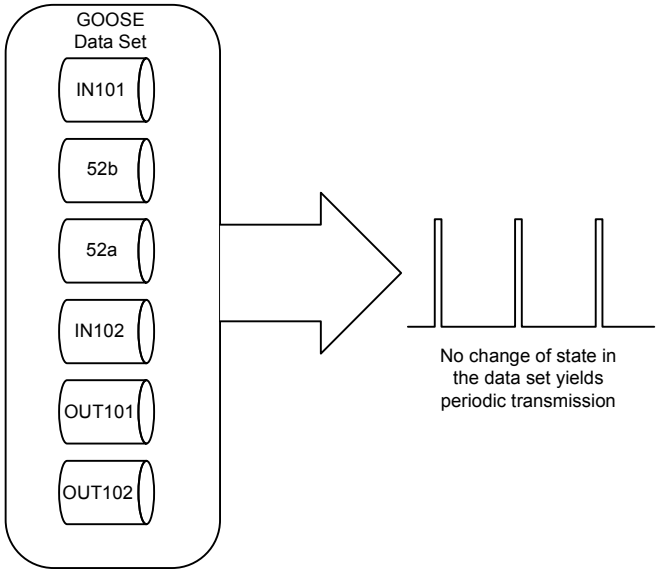

Fig. 6 GOOSE Message With Periodic Transmission

When the data set of the defined GOOSE message is static and none of the elements have changed state, the message is transmitted periodically over the network. Considering the relative size of the data set, this periodic transmission, generally on the order of once per second but sometimes user configurable, represents very low network traffic. To obtain protection-speed transmission times, the operational characteristic employed by GOOSE messaging triggers the broadcast of the message on a change of state of data within the data set. Accordingly, when an element within the data set changes state, a wave of message broadcasts is triggered in a very short time in order to ensure the message is received at the other end. Fig. 7 illustrates this mechanism.

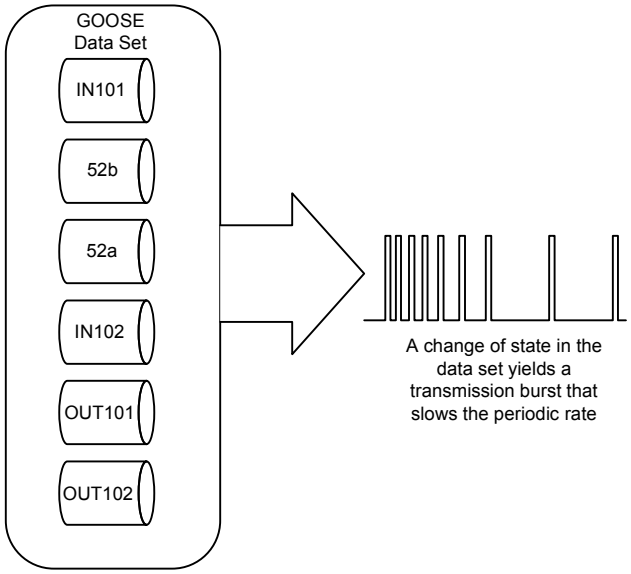

Fig. 7 GOOSE Message After a Change of State in the Data Set

Considering the operational characteristic of GOOSE messaging, the main question becomes: if several GOOSE messages are transmitted at the same time, how capable is the network of handling large bursts of information? In a worst-case scenario, a GOOSE message can be configured to carry as much information as can fit within a single Ethernet frame—roughly 1,500 bytes of data. A typical network installation uses switches capable of 100 Mbps throughput. A quick calculation, as shown in the following equations, yields roughly eight full-frame GOOSE messages occurring at the same instant of time, which saturates the capacity of a 100 Mbps port on a switch.

$$1{,}500 \text{ bytes} = 12{,}000 \text{ bits} \tag{1}$$

$$\frac{100 \text{ Mb}}{\text{second}} = \frac{100 \text{ kb}}{\text{millisecond}} \tag{2}$$

$$\frac{\dfrac{100 \text{ kb}}{\text{millisecond}}}{\dfrac{12 \text{ kb}}{\text{GOOSE message}}} \approx \frac{8 \text{ GOOSE messages}}{\text{millisecond}} \tag{3}$$

While these equations represent a worst-case scenario, it is easy to see how a flood of GOOSE traffic could impact the performance of the network. The GOOSE messages are broadcast most frequently upon a change of state of the data within the data set. It is reasonable to assume that the changes of state occur because of an event on the power system. Consequentially, the triggering of GOOSE traffic on the network may come at a time when reliable communication is most wanted, yet the very nature of GOOSE messaging may decrease that desired reliability. While this example may cast a dim light on the prospects of utilizing GOOSE effectively for network protection purposes, it is not meant to dissuade the use of GOOSE, only to make the readers aware that great care should be taken when designing the network and crafting the GOOSE data sets. There are simple tools built into the IEC 61850 standard that can aid in ensuring that the network switches process the messages in a priority-based fashion, as well as allowing the network to be segregated virtually.

IEC 61850 GOOSE communication allows for messaging to be segregated into virtual local-area networks (VLANs), as well as have a prioritization assigned to each message. VLANs allow network switches to block certain messages from going to certain ports within the device. When implemented properly, VLANs minimize the amount of traffic each IED receives and therefore reduce the amount of data the IED must process from the network. In circumstances where large amounts of data storm through the network, VLANs ensure that an IED does not have to process data unintended for the device.

In addition, data can also be prioritized when queued in a network switch. The priority assigned to a GOOSE message determines the sequence in which the message is processed.

Given the possibility of large amounts of GOOSE traffic converging on the system all at once (a bus fault, for example, could trigger many relays), it is important to consider such tools when using IEC 61850 GOOSE.

B. Note on IEC 61850 GOOSE and Performance Standards

It is worth noting that IEC 61850 classifies GOOSE performance into classes based on communications latency performance between a sending and receiving device. The standard identifies multiple classes and types, the most rigorous being Type 1A, Class P2/3. In order for a device to meet Class P2/3 latency standards, the latency between two devices must be 3 milliseconds or less. The important thing for the reader to understand is related to the data observed in [1]. As referenced, the performance of the network decreases when the network is more burdened. Therefore, two devices communicating to each other under low-burden networks may meet the Class P2/3 classification; however, if the network traffic is increased substantially, the same two devices may no longer meet the Class P2/3 classification. The important fact to absorb from this example is that the performance classification is only partially dependent on the actual device but also partially dependent upon the network congestion at any given moment of time. Consequentially, a device classified as meeting Class P2/3 criteria is not guaranteed to

meet the criteria because of its dependence on network conditions.

VII. PROTECTIVE RELAY NETWORK PERFORMANCE

While the authors were unable to perform testing related to the performance of IEDs under various network loading conditions because of time constraints, the topic is very important and should be considered. Testing done in [1] indicates that the performance of the protective relay may be greatly impacted by the amount of traffic on the network. Reference [1] found that while all relays tested performed well under steady-state or quiescent conditions, when the IEDs were burdened with higher levels of network traffic, the performance of every relay suffered to some extent, and some suffered more than others.

The important idea to take away from such testing is that network standards have not been developed to test the performance of devices under extreme network conditions. While IEDs undergo extensive environmental testing to validate performance under extreme temperature, vibration, or shock conditions, no such standardized test exists to quantify the performance of the devices under extreme network conditions.

Given the relative importance of testing products under extreme conditions, it seems reasonable to expect system testing under extreme conditions. End users may consider requiring such network tests in order to generate added confidence in their system.

VIII. TIME SYNCHRONIZATION OVER ETHERNET

Time synchronization has proven its worth throughout the years. However, the most accurate method of device synchronization remains a hard-wired signal sent directly to the device, such as an IRIG-B signal. Notwithstanding, there are network-based time-synchronization options available in NTP and SNTP. SNTP is the chosen method of time synchronization within the IEC 61850 standard. While theoretically SNTP can have submillisecond time accuracy, [2] acknowledges that practical implementations can expect accuracy to within 100 milliseconds. While this might be suitable in some applications, protective relaying time-stamping during system events requires accuracies of less than 1 millisecond.

However, Precision Time Protocol (PTP) promises to be a great improvement over past network time-synchronization methods and claims submicrosecond accuracies. PTP is being developed through IEEE 1588-2008 and will be available to end users in the future.

IX. OTHER CONSIDERATIONS

Although the main intent of this paper is to address the technical issues of latency, redundancy, and failover recovery, other outside requirements can affect system performance and functionality if not addressed and clearly understood. One key challenge for both power system engineers and PMLS system integrators is meeting the network security and information protection compliance requirements for today's systems. With the increasing need for secure networks and information protection, most companies today have information technology (IT) and/or IP policies in place that require stringent rules and procedures to ensure that compliance is met before systems are placed online. These methods are usually broken into two categories: approved IT system hardware and technical controls.

These policies usually require that all PMLS network hardware be reviewed and approved by the company IT department. The network bill of materials is usually submitted for approval by the project integrator. This review process can take some time, usually with the project team explaining the different types of equipment that are used in the integrated substations.

After the bill of materials is approved, the next step is implementing the technical controls on the network. Technical controls can consist of many processes, depending on individual policies. Some key examples that usually come into play are: secure logins, antivirus programs, standardized computer images, firewalls, network compliance monitoring programs, secure IP addresses, network switch controls, and so on. Each of these technical controls is clearly intended to provide methods of ensuring network security and compliance. Although these technical controls do a good job of ensuring network security, they can also have unintended consequences if system performance and functionality are not checked and verified.

It is clear to the authors that these hardware and technical control policies have been written from a corporate business network thought process. These policies are not necessarily a clean fit for process control network systems, but these systems are still required to meet the obligation. These policies work well for business environments, where downtime and timing issues are not as noticeable as with process control systems, where uptime and reliability are the key focus business drivers.

It is the authors' recommendation, based on previous experiences, that early dialogue and engagement between company IT personnel and the project team are essential to ensure that both IT and/or IP compliance requirements are met along with the intended system functionality. Early engagement provides the opportunity to clearly understand all requirements while providing enough time to test and document the approval process. It is better to engage early and receive approval and/or any needed waivers from IT rather than learn later that the system does not meet compliance requirements, risking startup delays.

X. CONCLUSION

The authors recognize the importance of network communication for modern power systems and the impact it has on the advancement of reliable and secure power system operation. In order to use this communication to its fullest and most useful potential, it is important for the implementers of such technology to understand weaknesses in the technology and design systems that address such issues. Modern network communication provides some tremendous advantages over older, more traditional technologies, and it represents a tool that engineers can use in various situations

for various applications. However, engineers must be sure that they are using the proper tool for the job, and some problems may best be solved using more traditional methods. It is up to the engineers to educate themselves to make the best decision.

The findings of this paper make it clear that engineers should be aware of several issues when deciding on using network communication in the substation, especially when performing protection communication over the network. The most important issues include the following:

- System end-to-end performance should be considered when judging the robustness of a network, yet the current network communications standards do not address such performance.
- Network traffic has an impact on the performance of the system and should be considered in the network design.
- When using IEC 61850 for protection communication, VLANs and prioritization are strongly recommended.
- Network-based time synchronization may not meet system performance requirements.
- Interfacing with end-user IT departments should be done early and often, with explicit declarations that standard IT-grade network equipment may not be sufficient.

XI. REFERENCES

[1] K. Leggett, R. Moxley, and D. Dolezilek, "Station Device and Network Communications Performance During System Stress Conditions," proceedings of the 1st Annual Protection, Automation and Control World Conference, Dublin, Ireland, June 2010.

[2] Cisco Systems Document ID: 19643, "Network Time Protocol: Best Practices White Paper," December 2008. Available: http://www.cisco.com/application/pdf/paws/19643/ntpm.pdf.

XII. VITAE

Nicholas Seeley graduated from the University of Akron in 2002 with a BS in electrical engineering. After graduation, he began working at American Electric Power in Columbus, Ohio, for the station projects engineering group, where he focused on substation design work. In June 2004, he was hired at Schweitzer Engineering Laboratories, Inc. in the engineering services division as an automation engineer involved in the development, design, implementation, and commissioning of numerous automation-based projects specifically geared towards power management solutions. He currently works as a lead power engineer in the research and development division.

Kurt Concienne received his undergraduate degree in Electrical Engineering from Lamar University in Beaumont, Texas, in 1989. He began working as an instrument and electrical engineer for Chevron at their Port Arthur, Texas, refinery. He then was the facility electrical engineer at the Chevron Chemical Plant in Baytown, Texas, working in maintenance and engineering and progressing to plant electrical section supervisor. Currently, he is working in Chevron's Engineering Technology Center as an electrical engineering power systems specialist.

Previously presented at the 58th Annual Petroleum and Chemical Industry Conference, Toronto, Canada, September 2011.

20110427 • TP6471