Choose Simplicity for a Better Digital Substation Design

Greg Rzepka, Scott Wenke, and Sarah Walling
Schweitzer Engineering Laboratories, Inc.

This paper was presented at the 70th Annual Conference for Protective Relay Engineers, College Station, Texas, April 3–6, 2017, and can be accessed at: https://doi.org/10.1109/CPRE.2017.8090016.

Choose Simplicity for a Better Digital Substation Design

Greg Rzepka, Scott Wenke, and Sarah Walling, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Digital substation process bus-based solutions provide multiple benefits that can be used to improve personnel safety around high-voltage equipment, reduce electromagnetic interference challenges, and decrease substation construction costs. To realize these benefits, special attention must be given to the selection of the fiber-optic network architecture used for a process bus implementation. Network architecture directly affects the amount of engineering required to design, commission, and maintain a digital substation.**

This paper outlines the benefits of using digital communications between primary substation equipment and the protective relays installed in the control house. The authors discuss different communications network architectures and their associated protocols and then introduce a point-to-point, EtherCAT®-based digital substation technology. Simple and secure, this new solution is easy to implement, eliminates the need for external time synchronization, requires no network engineering, and is easily scalable to support traveling-wave protection and the megahertz sampling requirements of future digital substations.

I. INTRODUCTION

The electric power system is one of the largest and most complex systems ever built. Continuous advancements in protection and control system technology implemented during the past century have resulted in significant improvements in power system dependability, reliability, and security while decreasing overall system cost.

The introduction of the microprocessor-based relay in the 1980s was a cornerstone for the gradual digitization of substations, offering unprecedented levels of data collection and communications as well as increased system integration. Throughout this transformation, one crucial part of the system remained virtually untouched by the progress—the substation yard copper wiring that connects instrument transformers and high-voltage apparatus control circuits to protective relays.

Technology advancements significantly reduced the amount of energy instrument transformers need to supply, resulting in a large mismatch between the instrument transformer drive capability (5 A, 115 V, 1 to 2 kVA) and the digital relays whose needs are limited to information about primary quantities. Modern digital relay inputs typically consume less than 0.5 VA per circuit. In effect, this means most of the instrument transformer energy is expended driving the unnecessary copper wiring loads.

The 5 A nominal current standard popular in the U.S. requires high cross-section conductors, resulting in expensive copper wiring. In a large substation, copper costs can easily exceed the cost of the protection and control equipment connected to it. This situation can be remedied by using optical fibers to bridge the gap between the substation yard and the microprocessor-based protective relays located in the control house. Data are transmitted digitally with multiple signals multiplexed onto a common fiber-optic communications medium. This approach is referred to as a digital substation solution.

Recent changes in substation architecture occurred with the introduction of the so-called station bus. A station bus is an Ethernet-based communications bus that allows all connected intelligent electronic devices (IEDs) to exchange digital data (see Fig. 1). This approach greatly simplifies copper wiring between IEDs, but it introduces complexity in the configuration and testing of the communications infrastructure.

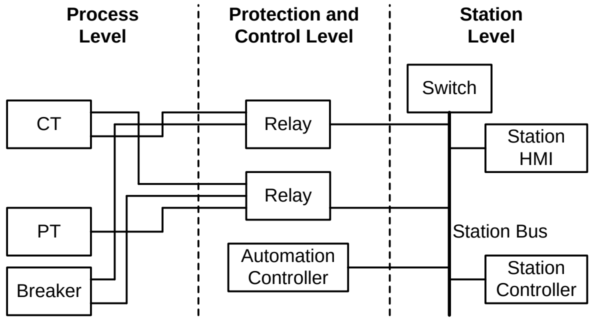

Fig. 1. Substation configuration with a station bus

Going forward, the industry is evaluating ways to replace wiring at the process level with fiber-optic connections. This can be achieved by digitizing analog currents or voltages at the measurement point and sending digital samples using communications protocols (see Fig. 2). Remote data acquisition devices used for digitization are called standalone merging units, or simply merging units (MUs).

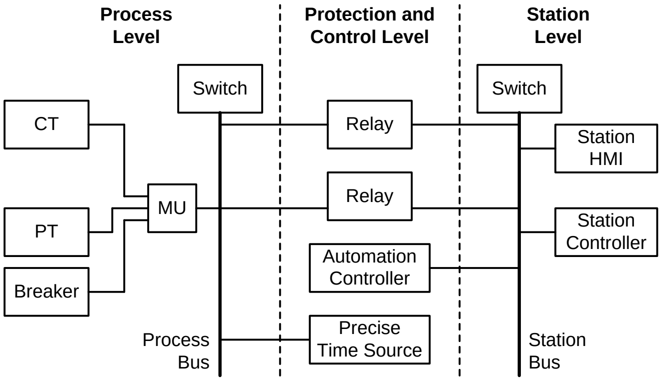

Fig. 2. Substation configuration with a station bus and process bus

This paper discusses how the proper application of a digital substation solution increases personnel safety, improves reliability and cybersecurity, and helps utilities save on construction and life-cycle expenses. The paper also explains two methods used to implement a process bus—a switched network and a point-to-point network—and compares various aspects of these architectures. Finally, this paper introduces a new point-to-point, EtherCAT®-based digital substation solution that is robust, simple to use, and secure.

II. Digital Substation Motivations

Motivations for using a digital substation technology include lower costs, improvements to personnel safety, and system reliability gains. When evaluating the potential advantages of using a digital substation, it is important to weigh the aforementioned benefits against the expected increases to engineering complexity and the risks associated with applying new technologies. When not fully understood, complexity can affect reliability and maintainability as well as significantly increase the total cost of ownership. The risk of implementing any new technology should be considered, especially the maturity of said technology or associated standard, the availability of technical support, and the workforce training requirements.

A. Safety

One benefit of separating the instrument transformers from the relay is that it reduces the risks to personnel safety, such as that of an exposed potential transformer (PT) or open current transformer (CT) connection. Because the PT or CT connection is now reduced to several meters and located in the yard, the control house environment where typical day-to-day operations occur becomes much safer overall and the accident risk is significantly lower. Shorter yard wiring also reduces the exposure to electromagnetic stresses presented to both the MUs and the digital relay equipment. Properly designed and grounded systems can increase overall substation resilience and improve survivability against high-energy electromagnetic pulse threats.

Another benefit of not having high-energy cables in the control house is that replacing devices or adding new devices can be performed quickly. Typically, an established set of procedures ensures that no one is working on an energized circuit and that there is a safe work environment. However, because the only signals going to the relay are fiber-optic connections, many of the processes can be sped up or obsoleted altogether.

B. Reliability

Solutions that use fiber are proven to be more reliable than those that use copper [1]. Fewer physical routing paths and connections are needed overall, which decreases the likelihood of wiring errors and missed connections. This also reduces the chance of inadvertent misoperations.

With analog measurements in digital form, modern microprocessor-based relays can self-monitor a larger portion of the secondary circuitry of protection systems. This allows for quicker identification of a problem and its removal before the wiring failure affects power system operations.

On the other hand, the increased number of IEDs required to implement a digital process bus impacts the overall mean time between failures (MTBF) rate for the system, which can be calculated as follows:

$$\frac{1}{\text{MTBF}} = \sum_{n=1}^{N} \frac{1}{\text{MTBF}_n} \qquad (1)$$

The typical MTBF rate of modern microprocessor-based relays is in the neighborhood of 300 years. For illustration purposes, assume the same MTBF for every device in the system. Under these assumptions, a system consisting of N devices with an MTBF of 300 years is proportional to a device's MTBF and inversely proportional to the number of devices, which can be written as follows:

$$\text{MTBF} = \frac{1}{N} 300 \qquad (2)$$

At a minimum, a digital substation requires two devices: an MU and a relay. Depending on the chosen architecture, additional devices may be required. For example, a switched network process bus requires an Ethernet switch, and protection systems using measurements from two different MUs require an additional high-accuracy time reference (e.g., a GPS clock). One can argue that additional redundancy can be built into the system to achieve the required levels of system availability. While true, the cost of maintaining a highly redundant system needs to be carefully considered.

Overall reliability goals for the system also need to be carefully considered. Increased reliability related to a lower number of termination points and greater ability to monitor digital links can be diminished by an increased number of electronic devices, dependency on communications infrastructure, and the need for a high-accuracy time reference.

C. Cost

Cost savings are another benefit to upgrading a protection and control system. Hundreds of feet of copper wires are replaced with just a few fiber-optic cables, reducing material costs and the physical footprint needed for trenches in the yard. Additional savings can be found by applying multifunction IEDs for protection, control, and monitoring purposes [2], though device cost varies based on the application.

Replacing copper with fiber also significantly reduces the time and labor required for system installation, documentation, commissioning, and maintenance. According to [3], 75 percent of the cost of installing a copper-based protection and control system in North America is related to labor. Traditional copper substations require thousands of individual connections that must be terminated one by one by skilled personnel, while a modern digital substation solution only requires a few fiber connections and moves wire termination work to the equipment manufacturer [4] [5]. It is easy to envision a future in which the MU is already installed in the breaker or IT cabinet by the device manufacturer, potentially eliminating the field wiring process and replacing it with computer-based signal mapping and path configuration.

Comparing traditional protection systems and new process bus technologies using an example system provides insight as to where cost savings can be realized as well as where costs can increase. For this discussion, a single line relay serves as the example system. Table I outlines the costs associated with activities required for developing a traditional protection system, a process bus-based protection system that uses a switched network, and a process bus-based protection system that uses a point-to-point architecture. An explanation of each activity and its costs is then provided. Note that for any given installation, the savings and costs change depending on the number of devices needed, the complexity of the system, and so on.

TABLE I
COST COMPARISON FOR CONVENTIONAL AND PROCESS BUS INSTALLATIONS

Activity	**Traditional Protection System**	**Process Bus Systems**	
		Switched Network	**Point-to-Point Network**
Engineering labor	$26K	$27K	$21K
Electronic devices	$7K	$11K	$10.3K
Copper	$6.6K	$1.2K	$1.2K
Fiber	NA	$2K	$2K
Other labor	$3.8K	$2.1K	$1.7K
Total cost	$43.4K	$43.3K	$36.2K

Engineering labor, which includes protection design, protection engineering, SCADA work, and drafting costs, makes up the bulk of expenses. Estimated costs for the three network types are based on standard rates from an engineering consulting firm along with estimates of how long the work would take for the various technologies. Labor costs increase for a switched network because while the labor for drafting and wiring diagrams is decreased, additional engineering and documentation work is needed to configure the network. Ensuring that the network is well documented for factory acceptance testing and for future troubleshooting can add up to an additional thousand dollars. It is worth noting that for large projects, engineering labor costs can be spread across several protection schemes that use the same network, resulting in a lower cost per application than what is described in this subsection. For a point-to-point network, savings are found through lower drafting and wiring costs. Moreover, because the network does not need documentation beyond the physical connections, similar to a conventional system, there are additional labor savings.

The cost of electronic devices varies for a traditional application, a switched network topology, and a point-to-point topology. Therefore, assume that the IEDs for all three technologies are similar in cost. Note that published device prices from a single manufacturer were used to limit the scope of the example and to be consistent when comparing network types. The additional costs for the process bus solutions come from additional equipment. The switched network also requires an MU, a switch, and a clock. The switch and clock can be used for other protection applications at the same site, so a discounting factor was applied to the costs shown in Table I. In the point-to-point network, the only additional device cost is the MU in the field.

The copper and fiber material costs can be grouped when talking about these solutions because one goal of a process bus solution is to replace many copper cables with a single fiber cable to save money. For the single line relay example, copper cables are assumed to include power, CT/PT connections, and control connections. Copper prices can be volatile; however, the example assumes a price of $3 per foot. For a conventional system, copper is used for all connections, so there is no fiber cost. For the process bus solutions, the costs are identical, with the copper costs primarily coming from copper runs to apply power to the MU or other equipment or for energizing contacts in the yard. While not outlined in Table I, other material expenses such as panels and control houses are largely the same no matter which solution is chosen.

The activity labeled as other labor includes most of the field work required for the installation and wiring of all the equipment. It is assumed that the MU is installed by the primary equipment manufacturer and that it comes installed and wired when ordered with equipment such as a breaker cabinet. In this case, the only connections the end user needs to make are the fiber and power connections. Having fewer connections and having to do less cable trench work results in significant labor savings. The increased labor costs for the switched network solution stem from the fact that additional connections need to be made and managed with the switch in the substation. Skilled personnel and network engineering tools must be used because the connections need to be precise in order for the solution to work with expected performance.

Considering all of these activities, the total cost of implementing a switched network is nearly the same as that of a conventional copper-based network. However, users can potentially see a 17 percent cost savings by implementing a point-to-point network.

III. PROCESS BUS ARCHITECTURES

When talking about a digital substation, it is important to make a distinction between and consider the different requirements for the process bus and the station bus. The station bus allows the exchange of information between IEDs in the control house, while the process bus is used to communicate unprocessed system information such as raw voltages, currents, status signals, and decision signals (i.e., trip signals) [6]. Because of the nature of these signals and how they relate to a protection system, the requirements of each network can be very different when considering latency, jitter, availability, and data loss. For example, periodic retransmission of messages guaranteed by Generic Object-Oriented Substation Event (GOOSE) protocol helps to protect against single packet loss. Protocols such as Rapid Spanning Tree Protocol (RSTP) allow for dynamic network reconfiguration in case of a communications link failure. The healing times offered by a protocol such as RSTP can be on the order of tens of milliseconds. It should be noted that critical data (e.g., a breaker failure message) can be transported on the station bus. In that case, the station bus requirements are similar but much higher than the process bus requirements.

For the process bus, the requirements for healing and availability are much more stringent because the process bus now includes real-time data that are constantly streaming and critical for protection. This means that not only are the availability and healing requirements stricter, but that the amount of data is generally much higher. As an example, an IEC 61850-9-2LE stream with a set of currents and voltages typically requires on the order of 5 Mbps for a stream with data sampled at 4.8 kHz. For these example parameters, the MU sends out a packet approximately every 200 microseconds. The availability on a per-packet basis for the subscriber IED varies depending on the manufacturer's design, but it can be assumed that most allow for at least a single dropped packet without disabling the device. Either way, these network requirements mean that the typical healing times for an RSTP system (tens of milliseconds) are not sufficient to preserve protection system integrity when a communications system component failure occurs. Some possible measures to deal with these strict network requirements are discussed later in this paper.

A. Switched Network Architecture

A digital substation using a process bus can be deployed with a switched network architecture. In this case, the system consists of several devices: a unit near the primary equipment that samples analog data to publish to the network (i.e., an MU), switches in the network to route data, an IED to consume the sampled data, and a clock (see Fig. 3). Because of the variable nature of Ethernet and switched networks, the data from distributed data acquisition units (DAUs) must be synchronized and time-aligned. This is typically done through a GPS clock signal, such as a pulses per second signal, an IRIG-B signal, or a Precision Time Protocol (PTP) signal.

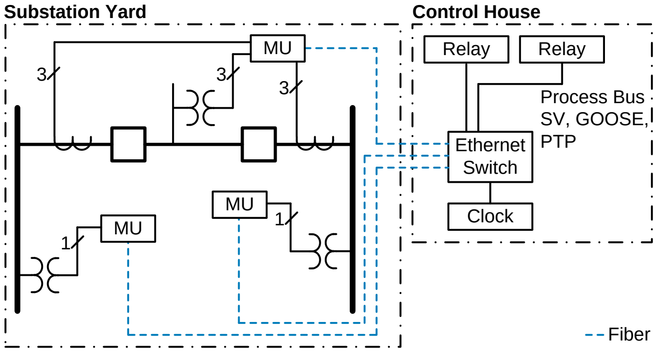

Fig. 3. Switched network architecture

One benefit of using the switched network method is that tools such as multicast packets can be used to deliver information from one producer to multiple consumers (e.g., multiple relay consumers subscribing to a single MU). However, because broadcast/multicast packets are sent to every device (even those that do not need or subscribe to the data), network management tools such as virtual local-area networks (VLANs) should be employed to ensure that a network does not become overloaded and adversely affect the operation of the protection system.

Another consideration and benefit of using a switched network is the ability to design a fault-tolerant network that can recover from any single failure (e.g., a fiber link or switch failure), ensuring that protection remains available. These design methods include RSTP to reconfigure networks automatically, Parallel Redundancy Protocol (PRP) to duplicate messages on separate networks, or software-defined networking (SDN) to pre-engineer failover paths in a deterministic way. While all of these methods provide healing for a network failure, each technology has a vastly different response time. RSTP heals in tens of milliseconds, SDN heals in less than a hundred microseconds, and PRP provides lossless failover over a duplicate network.

An important realization common to all network-based solutions is that as much as a protection engineer may want to abstract the communications network and think about it as a "cloud" capable of delivering messages where and when needed, it cannot be abstracted as long as the network is used to carry protection traffic. A protection engineer must be in a position to take full ownership of the network design and operation as well as have an intimate understanding of all failure modes and their interactions with the underlying protection scheme. Out of the technologies previously enumerated, SDN is the only technology offering the necessary level of control. Although very promising, ultimate success of this technology will be determined by a manufacturer's ability to offer a meaningful set of network configuration tools capable of empowering the protection engineer to take full control of the system design.

While the technologies described do provide redundancy for communications failures, it should be noted that no protection redundancy is provided unless a second protection device is also used in a scheme. This brings up further questions about the need for separate networks for completely independent and redundant protection; however, these are beyond the scope of this paper. These questions would need to be evaluated for each application to balance redundancy and resiliency with project costs.

B. Point-to-Point Architecture

A point-to-point digital substation architecture is just that—point to point between two devices. As shown in Fig. 4, data are directly sent from the MU in the yard to the protective relay in the control house. This removes a significant amount of complexity, including the switch, clock, and configuration and redundancy tools needed to serve a process bus application. Because the system is simplified, the network has lower latency and low jitter. It is also simpler to design because there is no need for tools such as VLANs or engineered SDN flows.

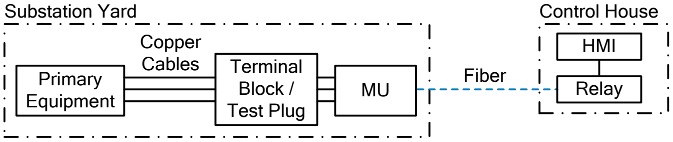

Fig. 4. Point-to-point architecture

IV. Comparison of Architectures

A. Equipment Connections

As previously mentioned, switched networks use four types of devices (MUs, relays, switches, and clocks), while point-to-point networks only use two types (MUs and relays). Both solutions use a terminal block at each site in the yard to provide a connection point for the wires from each external I/O point. The MU includes the terminal points for the I/O wires, so there are no immediate connections.

From there, the number of connections needed varies from one solution to the other. For a switched network (see Fig. 5), one fiber pair connects the MU to an Ethernet switch in the control house. The switch is then connected to a relay as well as to the time source. Because the switched network has more devices, it requires more fiber termination points. For a point-to-point network (see Fig. 6), the MU directly connects to a fiber-optic transceiver on the relay in the control house via a fiber pair. Even with multiple I/O connections, there are significantly fewer fiber terminal points required to connect a point-to-point system.

The number of connections required in a process bus network affects the reliability of the system [1]. Because a point-to-point system requires fewer connections, it is more reliable than a switched network.

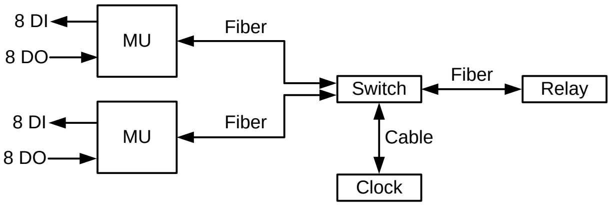

Fig. 5. Connections for a simplified switched network

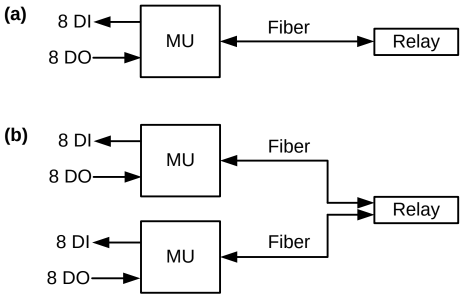

Fig. 6. Connections for a point-to-point system with a single I/O connection (a) and multiple I/O connections (b)

B. Time

Time synchronization is critical for a distributed system because of the need to align sampled analog signals for forms of protection, such as differential protection. If signals are not properly aligned, erroneous operations can occur.

When discussing time synchronization, there is an important distinction to make between absolute time and relative time as well as how they both work within process bus applications. For applications such as differential protection, the only requirement is relative time because the objective is to simply align local and/or remote samples.

Switched network and point-to-point architectures have very different challenges when it comes to maintaining relative time. For a switched network, the path that any one packet takes in the network can change every time. There are several reasons this can happen. If the network is heavily loaded, it can take longer to get the packet out of a switch buffer. If a link breaks, the physical path on which the packet is routed can change and take a longer or shorter amount of time to reach its destination. For these reasons, all devices in the system are generally synchronized to a common time source (often an absolute time source) so that all messages can be tagged and aligned later.

Time alignment is typically done with an IRIG-B or PTP reference signal that is delivered terrestrially using high-quality synchronous optical network (SONET) systems or derived from a source such as GPS or GLONASS. This means that all devices are synchronized to absolute time to account for any jitter the system may introduce. Alignment is performed in the end device (typically the subscriber) using sampled time stamps embedded in the stream that align all received data to a common time scale. Alignment fails if any of the data are asynchronous, giving preference to systems in which time is distributed along with the data (e.g., PTP) over implementations where time delivery system failure is independent from the communications network (e.g., IRIG-B).

In contrast, in a point-to-point network, time synchronization can be performed in a much simpler fashion. Because a point-to-point network has a fixed latency, with well controlled jitter generally in the order of 50 to 100 nanoseconds, relative time suitable for protection can be maintained by accounting for the fixed delay inherent to the connection. The benefit of this is that each connection can be measured directly by an IED and then calibrated on a per-link basis to create a relative time domain without the need for an absolute time signal. Another benefit is that because the methods to calculate the delay in the communications medium are fairly straightforward, they can be done automatically without user input. This makes a point-to-point architecture simpler to configure than a switched network. Because a point-to-point network does not require an absolute time signal, it has fewer elements, hence fewer points of failure.

C. Network Engineering

Using a switched Ethernet network for protection purposes generally requires network engineering expertise and tools. While it is possible to use an unmanaged network when the amount of data on a simple network is relatively small, this is not an advised practice considering the more stringent requirements needed to maintain process bus performance levels. Any unexpected consequence that arises from using the unmanaged network compromises the protection of the system. To this extent, tools for network engineering (e.g., VLANs and network priorities) are employed to ensure that network traffic arrives only at the desired location and that excess network traffic does not impact protection.

Using and understanding network engineering requires the configuration and coordination of several switches, which requires careful coordination and documentation. Using VLANs ensures that traffic is only routed to its intended recipient, with the switch only passing data with a specific VLAN tag to the specified ports. Giving each analog data stream and GOOSE message on the process bus a VLAN tag ensures that messages are only routed out of the ports that the engineer designates.

Another tool to route data accurately and quickly is to give the most important data the highest priority so that the switch can get the data through in a timely fashion. In the case of a process bus system, GOOSE messages that signal a trip and protection-critical analog data streams are examples of data that may be given higher priority due to their mission-critical nature.

While these tools are useful to ensure the integrity of a switched network process bus, they also require either the knowledge of both network engineering and protection requirements or, more likely, two or more individuals with these skill sets to coordinate and define application requirements for a successful deployment. This knowledge (or lack thereof) is something to consider when evaluating one of these projects.

Another challenge of a switched network architecture is testing the field installation against a number of abnormal conditions. Because switches typically use a spanning tree algorithm to heal from a point of failure, it is almost impossible to simulate every possible mode of failure and the network's healing behavior. Instead, a number of defined worst-case situations and typical modes of failure should be evaluated for a given installation to ensure that the network can handle enough modes of failure to be adequate for the application. Protocols such as PRP provide a way to design redundant networks by duplicating network infrastructure. This provides N-1 redundancy but at the cost of increased equipment and complexity.

Network-engineering a switched network is also challenging if a system is expanded after the initial installation. Being able to expand the network in the future using the same equipment to save on capital expenses is appealing. Many of the switches in the network may have additional ports that can be used to expand in a very cost-effective manner. However, the downside to this approach is that more data and new behaviors are introduced that need to be tested and evaluated to ensure that the new pieces of equipment work and that none of the previously validated applications are compromised. All of these challenges can be overcome, but they are considerations that should be evaluated for any process bus installation.

An external time reference is critical to distributed systems. Network-based process bus systems depend on the availability of a high-quality time distribution service. Contrary to the popular belief that time service can be implemented by connecting a few GPS clocks to the network backbone, reliable time distribution must be guaranteed and delivered by the network itself. Time distribution must be elevated to the level of a guaranteed network service to ensure (by design) that all devices that can communicate with each other also have the same notion of time. Time synchronization to an absolute time reference is less critical, but synchronization must be present as soon as multiple substations are connected together or to a common control center. Terrestrial time distribution is always preferred over wireless and GPS-based systems, which should be reserved for less critical local-area applications.

Substation network technologies are developing quickly, with a number of competing solutions already on the market. While the authors are optimistic that robust and mature switched network solutions will become available in the near future, equipment availability remains limited given the patchwork of mutually incompatible standards and systems resulting from the evolution of the technology.

A point-to-point system is easy to configure because of its simplistic design and direct connections. Wiring is standardized, is similar to traditional copper wiring, and is performed by equipment manufacturers. The number of individual cables is reduced because multiple circuits are served with a single MU and data are transmitted using a single fiber pair. A single fiber-optic cable can easily carry 20 to 30 fiber pairs, resulting in a significant reduction in the number of substation-hardened fiber-optic cables that traverse the yard. Circuits can be isolated and tested by simply unplugging a single fiber, requiring much less retraining and preventing accidental misoperations due to operator error. Personnel only have to validate that the connections are working, which is possible via visual inspection of the fiber port status LEDs.

Using standard industry fiber-optic cables makes termination and replacement easy and economical. Modern IEDs allow users to precommission a distributed system in a laboratory where it is easier to validate the protection system configuration. Detailed information about the configuration and topology can then be stored in the device's memory. When deploying the system in the substation, IEDs can verify that every module in the system exactly matches the precommissioned configuration. If end nodes are swapped or devices are wired in a different order, the issue can be quickly detected and resolved.

D. Cybersecurity

Process bus networks, which exist in the layer closest to the primary equipment, should be designed with cybersecurity best practices in mind. This is especially important because these networks are subject to North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP) compliance. NERC CIP provides high-level standards to help protect the physical security and cybersecurity of critical infrastructure, with CIP-005 to CIP-009 applicable to the process bus.

Switched networks and point-to-point networks are located and communicate within the physical security perimeter (PSP) and the electronic security perimeter (ESP). The PSP exists to prevent unauthorized physical access to an apparatus, and the ESP exists to protect communication within the network as well as with links outside of the PSP [7] [8]. Because of the location of switched and point-to-point networks, the need for cryptographically secure communication is reduced, and neither network has it built into its protocols [9].

Each device in a process bus is a potential target for attack because it provides a point of access to the network. While both switched and point-to-point networks use relays and MUs, the former also includes switches and clocks that must undergo their own security evaluations and then be managed via policies, network design, and engineering discipline to limit accessibility. Using static routes, whitelisting, and deep packet inspection that is available with the latest SDN technology helps with management, but configuration remains relatively complex. This added complexity can lead to misconfigurations and other human errors that increase security risks [9]. Given that a point-to-point architecture does not use a switch or clock and that it directly links an MU in the yard with a relay in the control house, there is no way to externally access the process bus [2]. The solution is therefore inherently cybersecure due to its simplicity.

Reducing or entirely eliminating access to the process bus helps avoid cyberthreats such as man-in-the-middle attacks. If access is possible, an attacker could disrupt the power system by altering data, modifying legitimate commands, or injecting malicious commands to trigger unwanted breaker operation. Because a point-to-point system is isolated and uses a direct connection for communication, the likelihood that the system will experience this type of attack is diminished.

E. Data Redundancy

Data redundancy is the ability to duplicate data and use them in case of lost data. This is important to consider when evaluating digital substation technology because with a process bus solution, more devices (and therefore more points of failure) are introduced. Data redundancy can offset reliability concerns by sending the information from the MUs out of multiple ports onto completely separate networks, like the method used in PRP. This means that the data are replicated; if they get lost on one network, they can still be consumed by the intended subscriber. Another way of looking at data redundancy is to have a single MU send data to multiple subscribers. If the data are lost or corrupted on the way to one subscriber, it is possible for the data to make it to the second subscriber and for that IED to use the data for protection. For either data redundancy method, care should be taken to ensure the data are duplicated and delivered in an expected way. If a subscribing IED receives too many unexpected packets it may, based on its design, disable to prevent any unwanted operations.

For a point-to-point architecture, data redundancy is more limited, but some benefit can still be realized. It is possible to maintain a point-to-multipoint connection so that data are duplicated to multiple devices but still maintain a point-to-point architecture between each pair of devices. This requires specialized hardware with multiple ports, but it can be used to strengthen the resiliency of the system against a single point of failure. The simplicity in a point-to-point system remains intact while gaining the benefit of data redundancy. In a point-to-point system, emphasis is given instead to functional protection system redundancy (e.g., Main I, Main II), which is well understood and widely accepted in the industry.

F. Ongoing Maintenance Issues

A big concern for users of these new process bus technologies is not only the initial installation but the ongoing maintenance necessary to ensure that the system stays running and is reliable throughout the life of the devices. Switched networks and point-to-point architectures have common maintenance concerns because the systems are now distributed instead of being in a central location, such as a control house. To test the analog-to-digital conversion circuits, controlled signals need to be injected at the site of conversion (i.e., at the MU) and verified through a metering check. This can be done most easily by checking a readout directly from the MU, if available, or by going through the subscribing equipment and accessing either the injection unit or subscriber with a remote connection. A third option would be to get the data that are going on the network or fiber directly and then decode the packets. This method can be sped up with a third-party tool, but typically it is not as simple and straightforward as the previously mentioned methods.

In addition to the common maintenance concerns, there are also concerns specific to switched networking that need to be considered when choosing a technology. One big concern is if any changes are made to the installation, such as installing new equipment or configuring devices to put additional traffic onto the network (e.g., a new GOOSE message). Any change to the network requires evaluation to determine additional network changes required, such as implementing a new VLAN for a new message. At a minimum, documentation must be updated. These updates and maintenance efforts require the review of a network engineer, preferably the engineer responsible for the initial design. This, in turn, brings about additional costs for any maintenance project that can make the overall cost of the switched network installation unsustainable.

Many of the described challenges can be eliminated through the use of a point-to-point fiber architecture. This is due in large part to the fact that a point-to-point architecture is deployed in installations today, just with copper instead of fiber. Many copper practices are analogous to practices that must be employed for a point-to-point fiber network. If there is physical damage to a system, the damage should be physically visible and can then be repaired or replaced as needed, similar to traditional installations. In addition, all fiber ports have LEDs to indicate their status, informing the user about link health, link activity, or other indications. These indications provide insight when trying to determine if the issue is with a port itself because the ports are usually running self-diagnostics. While there are a few more points of failure to examine in a point-to-point network, they can be physically seen, either through status LEDs or physical damage, which simplifies the maintenance and troubleshooting process.

V. New Point-to-Point EtherCAT Method

A point-to-point architecture simplifies many aspects of a digital substation solution, and additional benefits can be derived by using the EtherCAT protocol described in IEC 61158. EtherCAT was originally developed with a focus on short cycle times, low jitter, and accurate synchronization. Many of these same goals translate well into the requirements for a process bus solution. This section describes a new point-to-point EtherCAT-based digital substation solution.

In the proposed methodology for this solution, a single IED with several EtherCAT ports connects to several remote DAUs that have analog-to-digital conversion as well as digital input and output operation capability, all while maintaining a point-to-point architecture (see Fig. 7). Upon startup, each port explores what devices are connected in the network and creates a single, predefined packet that is passed between the IED and the DAUs. The packet is updated on the fly as it passes through each piece in the system. All of the data are then collected and aligned in the IED, where all protection and station bus functions are performed.

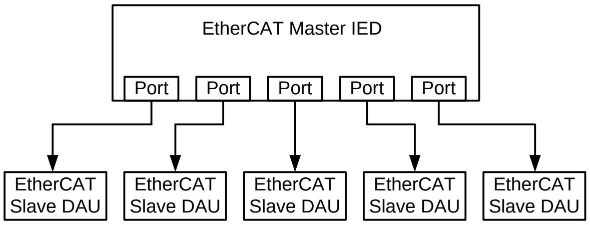

Fig. 7. An EtherCAT master IED uses multiple ports to establish point-to-point connections with multiple DAUs

Fig. 8 shows the operation of a traditional relay compared with an identically configured relay that uses EtherCAT technology. Overlapping COMTRADE captures show that the analog channels are properly aligned with the traditional relay due to IED compensation. The 2-millisecond delay in the overcurrent function operation is caused by the non-zero channel delay. The expected latency is less than 1.5 milliseconds; however, the device is processing the protection at 2-millisecond intervals.

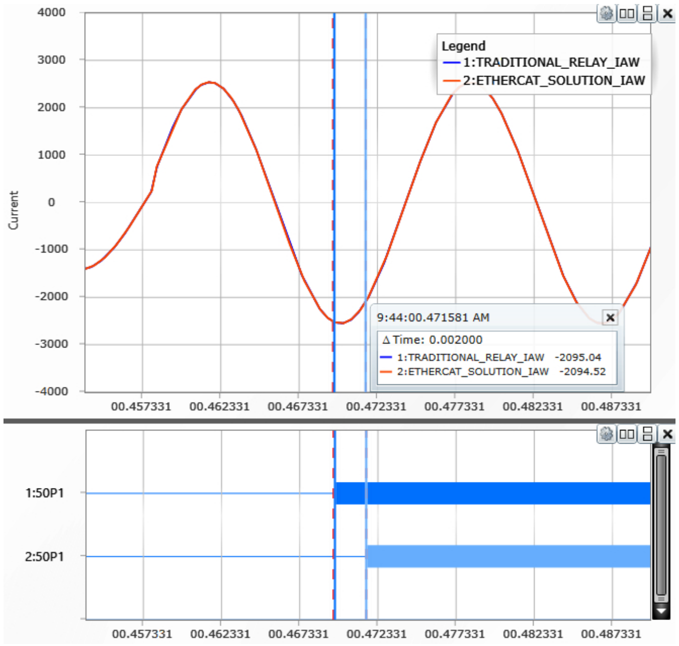

Fig. 8. Performance of a traditional relay and an EtherCAT-based relay

Relative time can be maintained by connecting the devices in a point-to-point fashion. In addition, the EtherCAT protocol has a mechanism for synchronizing multiple nodes so that they sample synchronously. By using a reference clock in the master and sending out synchronization messages, all slave nodes can be told to sample synchronously. Consequently, accuracy better than 1 microsecond (on the order of 50 to 100 nanoseconds) can be achieved.

Compared with standard Ethernet packets, an EtherCAT system makes more efficient use of available bandwidth. Because the connection is in a closed system with no switching, there is no need to build overhead into the packet for addressing and routing. In addition, the packet is constructed according to the actual devices found on the network, and therefore, there is no need to have placeholders in a packet for something that may not exist in that particular application. By having the packet predefined on startup, latency is reduced because the packet does not need to be parsed by any of the devices on the network and is written on the fly. This means the jitter time is solely dependent on how fast the packet can be read on the receive connection and routed to the transmit connection. This is a short enough time frame to essentially ignore the jitter in traditional protection applications.

Jitter and latency become more important as process bus applications such as traveling-wave protection are developed. This technology provides ultra-high-speed fault detection, often in the order of a couple of milliseconds. Traveling-wave protection requires megahertz sampling rates, or several orders of magnitude larger than the process bus solutions previously discussed in this paper. Given the amount of data transferred and the speed at which packets must be sent, any inefficiency in packet transmission is multiplied, making the EtherCAT method more suitable for high sampling rate applications like traveling-wave protection.

VI. Conclusion

Digital substation process bus-based solutions allow utilities to realize numerous benefits. Replacing copper with fiber removes dangerous voltages from the control house where personnel work, increasing safety. It also reduces the number of connections needed, thereby minimizing the likelihood of wiring errors and the subsequent rework required to address them. Because fiber's self-testing capabilities aid in early discovery of a corrupted transmission channel, either due to errors or a break in the cable, fiber is more reliable than copper.

Using copper in protection and control systems has a substantial effect on substation costs, going far beyond just the cost of cabling. By moving to a fiber-based solution, utilities reduce their material expenses and the labor costs related to designing, installing, commissioning, and documenting the system.

The switched network model and the point-to-point methodology are both suitable digital substation process bus-based solutions. However, a point-to-point system is simpler to engineer, deploy, and maintain because it does not require Ethernet communications infrastructure or an external, high-accuracy time source. Given the changing demographics of the workforce and the limitations of skill sets when it comes to network engineering, a point-to-point system is recommended when updating to a fiber-based solution.

The simplest point-to-point solution available is the new EtherCAT-based method described in this paper. It offers simplicity and security while solving many of the issues encountered with Ethernet-based solutions. First, EtherCAT maintains relative time between the relay and the DAU. The system does not rely on an external time signal for protection, and in a distributed system, the DAUs all sample synchronously with each other. Second, the solution provides low latency and low jitter because there is a direct connection for the predefined packet to travel between the publisher and subscriber. Third, it is scalable to accommodate future digital substation requirements, such as traveling-wave protection and megahertz sampling. And finally, it is easy to implement with no network engineering necessary.

VII. References

[1] G. W. Scheer and R. E. Moxley, "Digital Communications Improve Contact I/O Reliability," proceedings of the 7th Annual Western Power Delivery Automation Conference, Spokane, WA, May 2005.

[2] M. Adamiak, B. Kasztenny, J. Mazereeuw, D. McGinn, and S. Hodder, "Considerations for IEC 61850 Process Bus Deployment in Real-World Protection and Control Systems: A Business Analysis," proceedings of the 42nd CIGRE Session, Paris, France, August 2008.

[3] R. Hunt, "Process Bus: A Practical Approach," *PACWorld Magazine*, Spring 2009, pp. 54-59.

[4] J. M. Byerly, "AEP and Process Bus: Balancing Business Goals and Choosing the Right Technical Solution," proceedings of the 42nd Annual Western Protective Relay Conference, Spokane, WA, October 2015.

[5] B. Kasztenny, J. Mazereeuw, S. Hodder, and R. Hunt, "Business Considerations in the Design of Next Generation Protection and Control Systems," *Protection & Control Journal*, October 2008, pp. 15-18.

[6] E. Udren, S. Kunsman, and D. Dolezilek, "Significant Substation Communication Standardization Developments," proceedings of the 2nd Annual Western Power Delivery Automation Conference, Spokane, WA, April 2000.

[7] NERC Standard CIP-005-1 – Cyber Security – Electronic Security Perimeter(s). Available: http://www.nerc.com.

[8] NERC Standard CIP-006-1 – Cyber Security – Physical Security of BES Cyber Systems. Available: http://www.nerc.com.

[9] J. Casebolt, "Security Through Simplicity in Digital Substations," November 2016. Available: https://selinc.com.

VIII. Biographies

Greg Rzepka received his M.S. degree in electrical engineering from the Silesian University of Technology and Ph.D. in electrical engineering from the Missouri University of Science and Technology. Prior to joining Schweitzer Engineering Laboratories, Inc. (SEL) in 2005, Dr. Rzepka worked at Transmission System Operator in Poland. At SEL, he leads the protection systems group in research and development and is responsible for transmission, substation, distribution, and industrial product lines. He is a member of IEEE and IEC TC57.

Scott Wenke received his B.S. degree in electrical engineering with a power emphasis from Washington State University. Prior to joining Schweitzer Engineering Laboratories, Inc. (SEL) in 2013, Scott worked at Itron. At SEL, he is a product manager in the power systems group of research and development and is responsible for transmission and substation product lines. He has been a member of IEEE since 2012.

Sarah Walling received her M.S. degree in organizational leadership from Gonzaga University, her M.A. degree in journalism from the University of Oregon, and her B.A. degree in English literature from Sonoma State University. Prior to joining Schweitzer Engineering Laboratories, Inc. (SEL) in 2011, she worked at Purcell Systems and at Idaho National Laboratory as a technical editor. At SEL, she is a senior marketing specialist assisting with the marketing efforts for generator, substation, transmission, distribution, and motor product lines.

Previously presented at the 2017 Texas A&M
Conference for Protective Relay Engineers.

20170213 • TP6777-01