

Merging SONET and Ethernet Communications for Power System Applications

Edmund O. Schweitzer, III, David Whitehead, Ken Fodero, and Paul Robertson,
Schweitzer Engineering Laboratories, Inc.
 2350 NE Hopkins Court, Pullman, WA 99163 USA, +1.509.332.1890

Abstract—The type and amount of information exchanged between intelligent electronic devices (IEDs) within a power system are increasing. Initially, communication in protection systems was the exchange of a few bits of information to indicate the direction of a fault or to signal a remote relay to trip or block. The communications media consisted of hard wire, pilot wire, and other similar communications methods. In modern power communications systems, the range of supported applications includes line current differential schemes, synchrophasor data collection, supervisory control and data acquisition (SCADA), engineering access, voice, surveillance, pilot protection, event report collection, and many other types of tasks. Each application data type has different latency and reliability requirements.

Modern power communications systems are dominated by time-division multiplexing (TDM) communication, such as synchronous optical network (SONET) and synchronous digital hierarchy (SDH), and packet-based communication, such as Ethernet. TDM communication provides deterministic data by dedicating bandwidth to each data service. Ethernet incorporates a bandwidth-sharing scheme that allows each service to use bandwidth when it is available.

TDM communication is ideal for real-time protection and control applications due to deterministic characteristics. Packet-based communication is ideal for transporting event reports and performing similar services because these services generally do not have deterministic requirements and require the transport of large amounts of data. However, with the availability of pre-engineered Ethernet packet navigation techniques, such as virtual local-area network (VLAN) segregation and message priority, system designers are attempting to develop near real-time communications schemes with Ethernet systems, such as IEC 61850.

Too often, TDM- and packet-based communications are perceived as mutually exclusive and competing technologies. However, TDM- and packet-based communications can be integrated to operate together in a way that leverages the benefits of each technology.

This paper describes the benefits and shortcomings of TDM- and packet-based communications, when each technology is appropriate for different types of services and/or data, and how combining both transport technologies provides a robust network system that meets the needs of both real-time protection engineers and corporate information technology (IT) professionals.

I. INTRODUCTION

A power utility network contains a diverse range of equipment, with devices and applications that need to communicate at a local level within substations and at a wide-area level between substations and remote sites.

Typical power utility network services include voice, teleprotection, telemetry, video, control and automation, email, and corporate local-area network (LAN) access. From a communications perspective, these applications have different requirements in terms of latency, bandwidth, security, and fault tolerance.

There is a clear trend within the industry to move many of these applications and services to Ethernet-based communications systems, with the goal of reducing capital costs and standardizing on common interfaces to simplify network design and move away from legacy equipment when implementing system upgrades.

This paper compares the performance of time-division multiplexing-based (TDM-based) and Internet Protocol-based (IP-based) communications systems and examines the characteristics of each system in terms of suitability for different applications.

II. COMMUNICATIONS SYSTEM ELEMENTS

Before discussing the specific details of different communications systems, it is helpful to start with an overview of the key elements of a communications system.

Local-area communication is concerned with the network connectivity between end-user devices and applications that are physically located together. A LAN requires that all connected end devices have a unique address to maintain communication across a shared access medium. LAN connectivity in a power system network covers a wide range of physical interfaces, protocols, and data rates. For example, the Ethernet connection to a video camera for site surveillance defines a physical interface (RJ-45 connector), protocol (IP), and data rate (image resolution and frame rate). Similarly, the teleprotection interface on a particular relay may require an EIA-232 interface that supports a serial peer-to-peer protocol and a data rate of 38,400 bps. In power utility communication, local-area communication encompasses a wide range of interface formats and protocols.

Wide-area transport is concerned with the communication of network traffic between distributed LANs and control centers. These wide-area networks (WANs) involve the transportation of large amounts of data and interconnect many sites that support a wide range of applications.

Linking the LANs and WANs is a function called multiplexing, which provides access for LAN traffic to the WAN network, as shown in Fig. 1. Multiplexing deals with

how local-area data are assigned access and bandwidth to wide-area transport resources and manages data routing between source and destination entities.




Fig. 1. Communications System Elements

III. TRANSPORT TECHNOLOGIES

Many transport technologies have evolved over the years to provide a solution for communicating data over a WAN. These technologies fall into one of two categories: TDM- or packet-based communication.

TDM divides the shared transport medium or channel into a series of time slots, each with a specified payload size. Each service that accesses the shared transport channel is allocated a series of regularly repeating time slots. The incoming data streams are partitioned or segmented into separate blocks and allocated to appropriate time slots on the shared channel, as shown in Fig. 2.




Fig. 2. TDM Communication

To use an analogy, TDM can be compared to a ski lift used to transport N groups of skiers up a mountain. Each group has its own lift line that is allocated Chair N (time slot) on the lift. It takes the same amount of time for every person at the front of a line to get to the top of the mountain. The transport time per person can be determined and is constant.

TDM is most commonly associated with the bulk transportation of 64 kbps digitally encoded voice circuits. Synchronous optical network (SONET) and synchronous digital hierarchy (SDH) are two examples of TDM systems initially developed in the early 1990s to carry digital voice traffic, but since then, they have evolved into versatile data communications solutions. These systems have been widely

deployed in telecom applications, including the electric power industry.

For TDM-based systems, the traffic is broken into 8-bit data bytes and sent at a fixed rate.

For packet-based communication, the traffic is broken into a number of bytes (64 to 1,500 for Ethernet) and sent whenever a new packet is filled. As each data service has information to send over the shared transport channel, the data are queued in a buffer and sent sequentially. The data are transmitted over the packet transport link in the order the data frames are received in the buffer, as shown in Fig. 3.




Fig. 3. Packet-Based Communication

Contention resolution methods deal with the challenge of having multiple packets arriving at the same time and trying to access the shared transport channel. In this situation, data build up rapidly in the buffer. If the system is heavily loaded with many applications trying to send large amounts of data, it is impossible to buffer all the data—so frames or packets are dropped. Higher-level protocol deals with the detection of lost frames and data retransmission requirements.

Packet-based communication is obviously different from the TDM-based system. Returning to the ski lift analogy, unlike the previously described TDM ski lift where the members of a group wait in their own queue line for their group's chair (each Chair N), this ski lift has a single queue for all groups—there is no determinism in terms of latency. A better way to state the effect on latency is that it varies depending on the traffic loading on the lift and the number of skiers arriving at the lift line at any given time.

To stretch the analogy a little further, we compare how the two approaches manage the allocation of resources in the following two scenarios:

1. One of the groups has skiers waiting.
2. One of the groups has a burst of new arrivals, whereas other groups have few people with a steady stream of new arrivals.

In the first scenario, we would find that in the TDM-based lift, the lift attendant would allow a group's chair to go empty, whereas in the packet-based lift, no chair would leave empty.

In the second scenario, if a packet-based lift system were used, the large group would fill the (single) lift queue, adding a long wait time to the other groups. In contrast, a TDM-based lift would ensure that each lift line group had equal access to the chairs by effectively reserving chairs for the specific use of each lift line.

These examples illustrate an important difference between TDM- and packet-based systems in how a resource is

allocated to data services accessing the shared transport channel. In TDM systems, the resource is preallocated to each service and reserved for the sole use of that application. If the service has no data to send, the shared resource is not reassigned to another service. Instead, the time slots are transmitted empty. This means latency and throughput remain deterministic and constant for each service, regardless of the traffic loading on the system. In contrast, a packet-based system does not preallocate a resource to each application. If one application service has no data to send, the shared resource is allocated to another service, and no packets are transmitted if there are no packets ready to send. However, if a single service has a large volume of data to send, it will dominate the available system bandwidth. There are mechanisms developed for packet-based systems, such as priority queuing, class of service (CoS), and rate limiting, to prevent unbalanced bandwidth allocation, but the underlying principle of packet-based systems is to allocate bandwidth as efficiently and as readily as possible to any requesting service.

In summary, packet-based systems are considered more efficient in terms of resource or bandwidth utilization when compared with TDM but, as a consequence, are nondeterministic in terms of latency and throughput. TDM provides deterministic latency and throughput for each service, regardless of traffic loading on the system.

Ethernet is probably the most well-known and widely deployed packet-based transport technology.

Table I shows a comparison of TDM-based systems (SONET/SDH) and packet-based systems (Ethernet). When comparing SONET and SDH systems to Ethernet systems, SONET and SDH have traditionally been recognized as having the following advantages over Ethernet:

- Determinism, in terms of latency and bandwidth utilization.
- Low latency.
- The ability to dedicate bandwidth per application service.
- In-band operation, administration, and maintenance (OAM).

TABLE I
COMPARISON OF TDM- AND PACKET-BASED SYSTEMS

	TDM (SONET/SDH)	Packet (Ethernet)
Latency	Low	High
Deterministic	Yes	No
Bandwidth	Dedicated	Shared
OAM	Yes	No

Network management and OAM capability are incorporated into SONET and SDH through the allocation of in-band overhead data fields. This gives the technology the ability to reliably support all management functions associated with running, maintaining, administering, and repairing the network without negatively impacting the performance of data services using the network. In particular, in-band OAM gives

SONET and SDH the ability to rapidly recover from communications path failures, regardless of network size.

Ethernet does not inherently support OAM; it has required additional protocol development to support these functions. These protocols access the shared transport channel in the same way as any other data service and are subject to the same variances in latency and lack of determinism. In contrast, Ethernet offers the following advantages over SONET and SDH:

- More efficient use of bandwidth for “bursty” traffic.
- Ubiquity of Ethernet as an interface.
- Ability to support multicast and broadcast traffic.

Ethernet with IP has become a convergence protocol for many applications over the past 10 years, supporting an ever-increasing range of diverse services and applications. The connectionless approach of Ethernet means that packets are individually routed across the network without the concept of establishing an end-to-end connection between applications. This enables Ethernet to more efficiently utilize bandwidth.

However, neither SONET and SDH nor Ethernet has remained the same since being introduced. SONET and SDH have evolved to provide support for running IP and Ethernet services over TDM. Similarly, Ethernet has evolved to support virtual local-area networks (VLANs), CoS, and circuit emulation services (CES) to reduce latency and support circuit-based services.

There has been a growing debate over the relative merits of packet-based systems versus TDM-based systems as more and more services and applications migrate toward Ethernet. The debate is particularly strong in the power utility industry due to the predominance of TDM systems, diversity of applications, age of equipment, and safety-critical aspects of the data services being run over the network.

The introduction of multiprotocol label switching (MPLS) has added a new dimension to the packet-versus-TDM debate. MPLS offers improved performance over IP and Ethernet by providing OAM mechanisms to enable faster network recovery after system element failures, bringing it closer to the recovery times of SONET and SDH.

In the following sections, we examine the performance attributes of each technology when applied to specific power utility applications and analyze which transport technology is best suited to each application. We also discuss the technology migration challenge and examine the argument that one technology should be used over another or that a blend of technologies can provide a better solution.

IV. POWER UTILITY COMMUNICATIONS APPLICATIONS

A typical power utility system contains a diverse range of applications and control systems that require some form of data or voice communication. These applications cover the following functions:

- Substation control.
 - Local and remote substation control.
 - Supervisory control and data acquisition (SCADA) management.

- Substation data analysis (event reports).
- Real-time protection and automation.
 - IEC 61850 Generic Object-Oriented Substation Event (GOOSE) and Sampled Values (SV).
 - Teleprotection.
 - Relay protection schemes.
- Metering and power quality monitoring.
- Physical security.
 - Video surveillance.
 - Proximity alarms.
 - Access control.
- Voice communication.
- Corporate LAN access.

Enterprise information technology (IT) systems could also be included in the central control centers, but that part of the communications network is outside the scope of this analysis. The focus of this paper is the operation and support function requirements of the power utility network.

The utility applications listed can be categorized into three classes: non-real-time data (including substation control and data analysis, metering and power quality monitoring, physical security, and email and corporate LAN access), voice communication, and real-time data for protection and automation.

Reference [1] provides an excellent summary of the communications requirements for each class of application and is the basis for the information summarized in Table II.

TABLE II
UTILITY COMMUNICATIONS REQUIREMENTS

	Non-Real-Time Data	Voice	Real-Time Data
Delay (latency tolerance)	High (>100 ms)	Moderate to low (<150 ms)	Very low (<10 ms)
Asymmetry	High (<500 ms)	Moderate (<150 ms)	Very low (<4 ms)

Non-real-time applications are very well suited to Ethernet and IP when their high tolerance to delay and channel asymmetry is considered. Traffic also tends to come in short bursts for most applications, with the exception of video. Data loss and error detection are performed by higher-layer protocols in Ethernet systems and usually require the retransmission of data packets. Non-real-time applications can tolerate delays incurred by the retransmission of packets without impacting the service they support.

The IEC 61850 standard provides a complete framework for substation control and automation and is based entirely on Ethernet. IEC 61850 includes both non-real-time and real-time protocols for the control of substation functions and has been successfully deployed by many utilities worldwide. GOOSE and SV protocols are both part of the IEC 61850 standard and fit into the real-time class of applications. These protocols are designed to operate within a substation LAN where network latencies are minimal. Within this environment, Ethernet can be engineered to provide acceptable performance for these classes of applications.

Voice services are another example where Ethernet has proven itself to be a technology capable of meeting another class of utility communications requirements. Traditionally the domain of TDM, more and more voice services are moving to Voice-over-IP (VoIP) solutions.

This leaves real-time teleprotection services as the one application class that requires low-latency, deterministic behavior. Unlike GOOSE and SV data that remain within a substation LAN, teleprotection information travels long distances between protective relays on power transmission and distribution systems. Teleprotection data from protective relays represent the most critical information transmitted across a power utility network. Teleprotection signals communicate trip signals, line current differential data, system stability information, synchrophasor data, remedial action and special protection scheme information, and so on to maintain a safe and secure power system state. They directly prevent severe damage and personal injury from faults in a high-voltage power system. The maximum operation time for teleprotection messages is 4 to 10 milliseconds. In addition, legacy differential protection channels must have minimal symmetric delays on the transmit and receive paths. A typical tolerance of up to 4 milliseconds of asymmetry is the specification for most modern relays. Teleprotection data rates are relatively low (38,400 bps), but information is being sent continuously between relays to communicate relay status and breaker open and close commands. There is a very low tolerance to bit errors and no tolerance to data loss or interruption in the communications path. False trips caused by problems in the data communications network are unacceptable. However, the toughest requirement is the ability to provide uninterrupted relay communication during a recovery in the communications network due to fiber break or component failure. This requires the use of redundant backup communications paths or network self-healing ring designs with less than 5-millisecond switching times.

Traditionally, redundant, analog audio tone; power line carrier; and/or TDM have been the predominant technologies used to meet the performance requirements of teleprotection.

In the following section, we look at actual performance figures from the latest TDM, Ethernet, and MPLS systems and compare their suitability for supporting power utility operation and support applications.

V. PERFORMANCE COMPARISONS

Many papers have been written discussing the performance of Ethernet transport solutions for power system applications, including protection schemes [1] [2] [3] [4] [5]. These documents compare the performance of TDM-based protection schemes with the latest generation of Ethernet-based schemes, which include Ethernet- or IP-based multiplexers and MPLS transport systems. Many of these documents compare the system performance of these new Ethernet-based schemes to the performance requirements defined in the SONET standard [3] [4] [5]. Most of these publications present the data in a way that shows favorable performance of the Ethernet-based schemes over traditional

TDM system performance. Unfortunately, these comparisons are not exactly apples to apples and can be misleading.

A. Ring Break Restoration Times

One common misrepresentation of ring restoration is the comparison of ring healing times. Some white papers compare MPLS network healing times to the 50-millisecond healing time stated in the SONET standard GR-253-CORE. A 50-millisecond restoration time is defined by the SONET standard; however, for real-time protection, 50 milliseconds was never widely accepted. Many documents published over the past 20 years discuss the need for network healing and resynchronization times for protective relay applications in the order of 10 milliseconds or better. These faster healing times are common and dominate in current power system protection networks [2]. Table III provides a comparison of the ring break healing times between standard telecom equipment and substation telecom equipment.

TABLE III
SWITCHING AND SYNCHRONIZATION DELAYS

Transport Level	Standard Telecom	Substation Telecom
SONET ring switch	50 ms	5 to 10 ms
MPLS ring switch	50 ms	50 ms
DS-1 reframe	50 ms	1 to 5 ms

Unfortunately, this is still not the worst-case scenario for these comparisons. Communications network equipment designed specifically for power system protection applications operates from the transport rate (SONET or T1) to the application rate, which is 64 kbps for line current differential or teleprotection. The less-than-10-millisecond healing time includes the availability of data at the application level. Many of the Ethernet system performance numbers include the availability of data from the transport layer and do not include the resynchronization times of the DS-1 multiplexer, which is required to convert this signal back to the 64 kbps synchronous signal used by the relay or teleprotection device. The additional DS-1 reframe time is also included in Table III. This represents an additional restoration delay. The advantage of using systems that operate from the application level (relay) to the transport level (interstation fiber) is that healing times are all-inclusive.

B. Network Latency

End-to-end network latencies are the accumulation of delays added as the signal crosses different parts of the network. TDM and Ethernet ring networks can be configured using several different methods. However, a comparison of the latencies and how the latencies are accumulated can be performed. This paper provides a comparison of the most common architectures used in MPLS, IP multiplex, T1, and SONET systems. The data used to create this comparison were taken from the best-case performance data for the Ethernet

systems and the best-case performance data for the TDM systems from publications [2] [3] [4] [5]. The data used are only for products designed for substation and protective relay applications. A comparison of the delays accumulated in the two most common TDM substation applications is shown in Fig. 4.

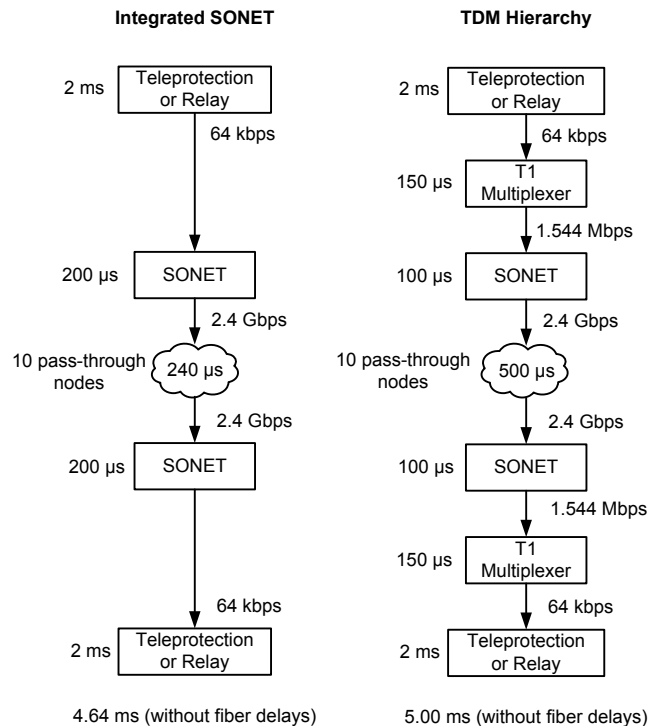


Fig. 4. Propagation Delay Buildup of SONET/SDH WAN Systems

TDM architectures, whether traditional buildouts or SONET multiplexers that operate to the application level, provide similar propagation delay performance. The example shown in Fig. 4 includes the time for a signal to pass through ten intermediate nodes. The same 4-millisecond back-to-back teleprotection or relay operate time was used for all comparisons.

Several Ethernet multiplexers are currently available that are designed for power system protection and control applications. These products use jitter buffers (which will be described in more detail later) and control how many DS-0 bytes are encapsulated per packet. These are the primary variables that exist to approach the propagation delay performance of the TDM-over-Ethernet conversion. Fig. 5 depicts two methods commonly used for power system protection and the performance results for each method. Note that the CES needed for accessing traditional interfaces usually offers the user the ability to change the samples-per-packet (S/P) size to trade bandwidth for improved propagation delay performance. Longer propagation delays are mainly attributed to the TDM-to-packet conversion, as shown in Fig. 5.

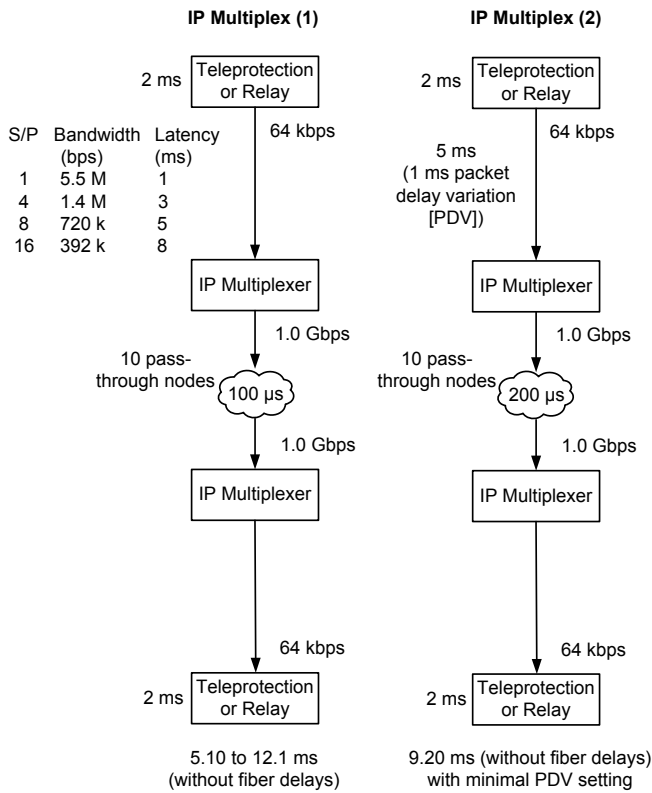


Fig. 5. Propagation Delay Buildup of IP Multiplexer Systems

Note that though MPLS does provide enhanced routing (compared with the traditional IP routing algorithms) with faster restoration times because of its ability to predetermine routes, it has the same latencies shown in Fig. 5.

From a power system perspective, why do we care about increasing overall fault-clearing times? The following excerpt from [6] provides this insight:

Power system faults and disturbances cause oscillations in the relative positions of machine rotors that result in power flow swings. The difference between a stable (return to a new equilibrium state) and unstable (loss of synchronism between groups of generators) swing is directly affected by the fault-clearing speed. Subcycle distance elements, along with the use of faster breakers, improve the likelihood of preserving power system stability during these conditions....

Subcycle distance elements also reduce the duration of through faults on transformers, which, in turn, reduces accumulated mechanical damage and extends transformer life.... [6]

C. Channel Asymmetry

Channel asymmetry is typically only a concern for line current differential schemes. New relays from all manufacturers can handle reasonable amounts of data channel asymmetry (4 milliseconds). All TDM systems designed for protective relay applications prevent asymmetry in the

transmission path by design. This is accomplished when transmit and receive paths for a circuit are not allowed to take different directions around the communications ring.

In Fig. 6, the blue path is desirable because it will have symmetrical delays and yield the shortest overall channel delays. The path shown in red will produce channel asymmetry and is not allowed through system programming. Once the paths for each direction of communication around the ring are set in a TDM system, the delay times and characteristics will be very consistent.

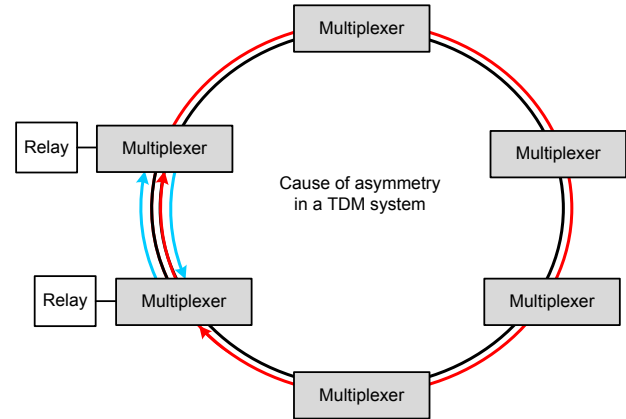


Fig. 6. Asymmetry Mechanism in a TDM System

In an Ethernet-based multiplexer scheme, channel asymmetry exists all the time. Unlike TDM systems, variable delays in Ethernet-based schemes are normal and “as designed.” All traffic entering an Ethernet network does so asynchronously, which means that packets from multiple ports may enter a switch simultaneously; the switch processes these packets in the best order it can. Priority queue settings can help manage these situations; however, if a packet is in the process of being sent, the next packet always has to wait.

D. Latency Details

At each egress switch port, a high-priority packet may have to wait for a maximum-length lower-priority packet to egress; a 1,518-byte packet takes 122 microseconds at 100 Mbps and 12 microseconds at 1 Gbps.

A potential 2-millisecond extra delay could, therefore, be incurred for a network path comprising 16 hops at 100 Mbps and 160 hops at 1 Gbps.

At each egress switch port, a high-priority packet may also have to wait for many other high-priority packets to egress; a 600-byte packet (typical for GOOSE) requires 48 microseconds at 100 Mbps and 4.8 microseconds at 1 Gbps.

A potential 2-millisecond extra delay could, therefore, be incurred for an event-triggered burst of 40 GOOSE packets at 100 Mbps and 400 packets at 1 Gbps.

Note that the latency of critical traffic can only be guaranteed if the nature (packet lengths and timings) of **all** traffic with the same or higher priority is well known and understood.

To avoid creating asymmetric delays, jitter buffers are used in Ethernet networks. A jitter buffer at each end of the line is

used to offset delay variation by queuing sent and received packets. The length of these queues is a compromise between limiting the effect on the overall propagation delay through the system and losing packets due to buffer overflow; this is a very difficult choice.

MPLS allows the user to fix the packet route through the network. This resolves the path-related timing issues of complex networks and limits excessive propagation delays due to signal routes (but, of course, it does not mitigate the queuing latency and buffer problems).

The asymmetry data shown in Table IV were collected over a 16-hour time period and taken from a line current differential relay operating over an Ethernet multiplexer system. This system was operating under ideal conditions, and the measurements were taken between adjacent nodes. The asymmetry histogram demonstrates that these asymmetry variations normally occur on Ethernet-based systems. The concern is what happens to these delays on nonadjacent nodes over a more heavily loaded system? The data in Table IV are from a single test site over a limited time; the data are not conclusive. Currently, the published data for channel asymmetry over these Ethernet systems is anecdotal at best.

TABLE IV
ASYMMETRY HISTOGRAM

Asymmetry (ms)	Percent Received
0.00 to 0.25	98.03
0.25 to 0.50	0.95
0.50 to 0.75	0.36
0.75 to 1.00	0.31
1.00 to 1.50	0.34
1.50 to 2.00	0.01

These asymmetries are introduced by the data entering the network asynchronously; some packets might be in contention at random intervals, creating occasional asymmetries.

These histogram data show that when using Ethernet as the transport technology for current differential relaying with TDM interfaces, analysis needs to be performed to further understand how these systems perform under all network conditions and what the maximum asymmetry delays that can occur are. New current differential relays have the ability to use the Global Positioning System (GPS) to synchronize the line current data and are highly recommended for operation over Ethernet networks. Current differential relaying applications on large, heavily loaded Ethernet networks require careful network design and studies of the maximum asymmetries that can occur during adverse system conditions. The bottom line is that all of the performance characteristics that make TDM ideal for current differential relaying are not the same for TDM-over-Ethernet multiplexing schemes.

VI. TDM AND ETHERNET

The best teleprotection performance over Ethernet is realized when Ethernet is the native protocol of the teleprotection device. The biggest performance hit for TDM-over-Ethernet systems is the process required to “packetize” the TDM data.

The IEC 61850 GOOSE message was designed to replace dc control wiring between devices. The performance requirement for GOOSE is an operation time of 4 milliseconds in the LAN. Many newer protective relays support GOOSE messaging. The main issue with using GOOSE as a teleprotection signal is that it is a Layer 2 broadcast message. This means that if GOOSE messages were used as the primary pilot communications signal (across the WAN), the network could become heavily loaded during a major power system disturbance.

The IEC 61850 GOOSE message includes VLAN tags, which enable message filtering at the trunk and port level of the network. This filtering allows the switch to block undesired messages from the devices connected to the network, but it does not help with data congestion at the transport level.

This is where using Ethernet-over-SONET (TDM) communication provides a performance advantage. The strength of a TDM transport is that bandwidth can be dedicated to applications, as shown in Fig. 3. When TDM bandwidth is provided for Ethernet, the attributes of a packet system are realized while inheriting all of the additional bandwidth segregation of a TDM transport system.

For example, the transport path through the WAN is predetermined and fixed, providing deterministic transport across the WAN for Ethernet communication.

Multiple TDM “pipes” can be created and used to carry Ethernet traffic throughout the network. This is the equivalent of having multiple, separate Ethernet WANs. An additional benefit of using TDM pipes is the ability to share GOOSE messages across the WAN while isolating them from the rest of the WAN traffic. Multiple pipes can be used to provide isolation between various services, such as IP phones, engineering access, GOOSE, IP surveillance cameras, or any service in which separation of the network traffic is desirable. Fig. 7 shows an example of two pipes with different bandwidths. Some products provide as many as 32 pipes for this purpose. These pipes provide complete separation of the bandwidth, as shown in Fig. 7. The bandwidth in the 150 Mbps pipe can be completely consumed with traffic from IP surveillance cameras and other security devices without affecting the performance of an interstation IEC 61850 GOOSE message on the 10 Mbps pipe. All TDM-based communications, such as current differential relaying and teleprotection, also coexist on this WAN system and perform with the same latencies that are expected and required for these systems.




Fig. 7. Ethernet Pipes Carried by TDM

Another benefit of Ethernet-over-TDM pipes is that healing times for fiber breaks in the WAN now occur at TDM speeds. This means that Ethernet traffic as well as TDM traffic will only be interrupted for typically less than 5 milliseconds for a fiber break in the system, regardless of the number of nodes in the system. Table V summarizes the preferred combinations of Ethernet and TDM teleprotection systems applied over Ethernet and TDM transport systems.

TABLE V
TRANSPORT PREFERENCE VERSUS TELEPROTECTION INTERFACE TYPE

Teleprotection	Transport	Preference
Ethernet	Ethernet	Preferred (LAN)
TDM	TDM	Preferred (WAN)
TDM	Ethernet	Not preferred (WAN)
Ethernet	TDM	Preferred (WAN)

VII. CONCLUSION

Both TDM- and packet-based communications technologies have unique advantages, which complement each other and allow the user to realize the best communication on a per-application basis. Applications that are point to point and channel-delay sensitive realize the best performance over TDM-based systems. However, applications that are not time-critical and require point-to-multipoint services, such as SCADA, engineering access, video, and (more recently) IP telephones, are most efficiently handled with packet-based systems.

One common comparison of TDM versus Ethernet systems is that Ethernet provides the most efficient use of bandwidth. This is only true for packet-based protocols transported over Ethernet networks. When TDM services are applied over these networks, this efficiency is lost. As much as 5.5 Mbps can be required to transport a 64 kbps TDM channel across an Ethernet network, as demonstrated in Fig. 5. However, when Ethernet is transported over TDM, bandwidth efficiency is maintained. This also provides the additional benefits of fast healing times and data segregation through the use of TDM pipes.

Overall, protection scheme operation times, which include protective relay detection time, teleprotection system time, and circuit breaker operate time, have been driven downward over the past decade. For protection at the extra-high-voltage

level, increasing the clearing times by 5 to 7 milliseconds is a step in the wrong direction.

VIII. REFERENCES

- [1] S. Ward and W. Higinbotham, "Network Errors and Their Influence on Current Differential Relaying," proceedings of the 64th Annual Conference for Protective Relay Engineers, College Station, TX, April 2011.
- [2] S. Ward, T. Dahlin, and B. Ince, "Pilot Protection Communication Channel Requirements," proceedings of the 57th Annual Conference for Protective Relay Engineers, College Station, TX, March 2004.
- [3] Alcatel-Lucent IP Division, "Teleprotection Over an IP/MPLS Network: Technical Validation," April 2011. Available: http://docbox.etsi.org/Workshop/2011/201104_SMARTGRIDS/03_COMMUNICATIONS/ALCATELLUCENT_Levrau_MPLSforTeleprotection.pdf.
- [4] RAD Data Communications, "Network Migration for Utilities – Teleprotection Over Packet."
- [5] J. Benckenstein, "Ethernet as a Channel for Protective Relaying," proceedings of the 64th Annual Conference for Protective Relay Engineers, College Station, TX, April 2011.
- [6] G. Benmouyal and K. Zimmerman, "Experience With Subcycle Operating Time Distance Elements in Transmission Line Digital Relays," proceedings of the 37th Annual Western Protective Relay Conference, Spokane, WA, October 2010.

IX. BIOGRAPHIES

Dr. Edmund O. Schweitzer, III is recognized as a pioneer in digital protection and holds the grade of Fellow of the IEEE, a title bestowed on less than one percent of IEEE members. In 2002, he was elected a member of the National Academy of Engineering. Dr. Schweitzer received his BSEE and MSEE from Purdue University, and his PhD from Washington State University. He served on the electrical engineering faculties of Ohio University and Washington State University, and in 1982 he founded Schweitzer Engineering Laboratories, Inc. (SEL) to develop and manufacture digital protective relays and related products and services. Today, SEL is an employee-owned company, which serves the electric power industry worldwide, and is certified to the international quality standard ISO-9001.

David Whitehead, P.E. is the vice president of research and development at Schweitzer Engineering Laboratories, Inc. (SEL). Prior to joining SEL, he worked for General Dynamics Electric Boat Division as a combat systems engineer. He received his BSEE from Washington State University in 1989, his MSEE from Rensselaer Polytechnic Institute in 1994, and is pursuing his PhD at the University of Idaho. He is a registered professional engineer in Washington and Maryland and a senior member of the IEEE. Mr. Whitehead holds several patents with several more pending. He has worked at SEL since 1994 as a hardware engineer, research engineer, and a chief engineer/assistant director and has been responsible for the design of advanced hardware, embedded firmware, and PC software.

Ken Fodero is a research and development manager for the communications product lines at Schweitzer Engineering Laboratories, Inc. (SEL). Before coming to SEL, he was a product manager at Pulsar Technology for four years in Coral Springs, Florida. Prior to Pulsar Technology, Ken worked at RFL Electronics for 15 years, and his last position there was director of product planning.

Paul Robertson is a product manager for the wide-area networking group at Schweitzer Engineering Laboratories, Inc. (SEL). He has over 15 years of experience developing and marketing products for the telecommunications industry, spanning cellular wireless and wire line communications systems. Paul worked in various technical and marketing roles for Motorola, Hewlett-Packard, and Agilent Technologies before joining SEL. He has a BEng in electrical and electronic engineering from Strathclyde University and an MBA from Edinburgh Business School.