Using Wavelength Division Multiplexing for Protection Applications

V. Skendzic and N. Fischer
Schweitzer Engineering Laboratories, Inc.

J. Sykes, D. Day, and K. Fennelly
Pacific Gas and Electric Company

Presented at the
CIGRE Study Committee B5 Colloquium
Tromsø, Norway
June 24–28, 2019

Previously presented at the
South East Asia Protection, Automation and Control Conference, March 2019

Originally presented at the
14th International Conference on Developments in
Power System Protection, March 2018

Study Committee B5 Colloquium
June 24-28, 2019
Tromsø, Norway

(Number)

Using Wavelength Division Multiplexing for Protection Applications

V. SKENDZIC*, N. FISCHER*, J. SYKES†, D. DAY†, and K. FENNELLY†
*Schweitzer Engineering Laboratories, Inc.
†Pacific Gas and Electric Company
USA

Summary

This paper reports on an experimental investigation that uses coarse or dense wavelength division multiplexing (CWDM, DWDM) for applications in high-speed traveling-wave protection. This paper documents the performance, opportunities, and pitfalls associated with this application and outlines practical strategies for the seamless integration of protection systems with the latest generation of optical transport network (OTN) technologies.

Keywords

Communications-based protection, wavelength division multiplexing, optical transport network.

1. Introduction

Power system protection typically operates autonomously, with communications-based schemes historically reserved only to protect the most critical transmission assets. Communications bandwidth and the cost associated with providing it has for many decades forced power system protection engineers to economize on communications resources.

Communications-based protection scheme deployments started with analog pilot wire, which sent current transformer secondary signals over a span of several kilometers. Pilot wire schemes evolved rapidly, spanning ever-increasing distances. Power line carrier technology, microwave links, and eventually time-division multiplexing (TDM) channels, such as synchronous optical network (SONET) and synchronous digital hierarchy (SDH), came next. They enabled the deployment of full segregated digital phase differential schemes (ANSI Device #87), as illustrated in Figure 1.

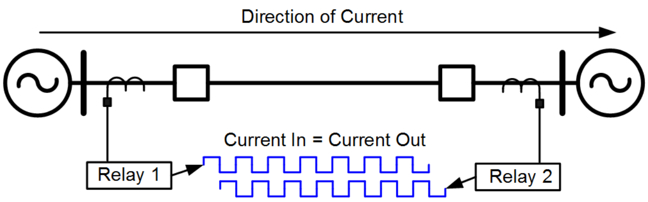

Figure 1. Differential line protection example (87L).

The Internet-driven data deluge followed, forcing communications systems to carry ever-increasing amounts of packet-based network traffic supplied by Ethernet and multiprotocol label switching (MPLS). Packet-based networks opened new interoperability challenges while simultaneously creating new opportunities to improve the overall quality of service [1].

The latest generation of ultra-high-speed (UHS) transmission line protective relays has recently raised the bar by introducing time-domain protection elements, traveling-wave fault location, high-resolution event recording, time-coherent MHz-level sampling, and current traveling-wave-based differential protection (TW87 element).

Without diving into a discussion about whether a dedicated TW87 fiber-optic channel is economically justified, it is sufficient to note that the sustained data bandwidth consumed by continuously transmitting 6 analog measurements (3 voltages and 3 currents) sampled one million times per second is on the order of 145 Mbps. Data are transported using a 1 Gbps Ethernet channel with small form-factor pluggable (SFP) transceivers, allowing the user to match the required transmission distance and control the associated light wavelength. As practicing engineers will immediately note, requesting a dedicated dark fiber channel for a single protection service is guaranteed to cause major interest in the communications department.

This paper (a shortened version of [2]) presents the results of a UHS protective relay test using a dedicated fiber-optic communications channel. The testing was conducted at the Pacific Gas and Electric (PG&E) High Performance Communications Technology Laboratory in San Ramon, California. The test was performed using the PG&E optical transport network (OTN) system, which carried the required relay payload with ease.

The paper also raises several questions about the design of future protection systems, revolving around the realization that modern fiber-optic communications systems have managed to exceed the communications bandwidth typically requested by present day protective relay designs by as much as 5 to 9 orders of magnitude.

2. Dedicated Communications Channel Requirements

A UHS transmission line relay TW87 communications link is a dedicated, private point-to-point fiber connection between two relays and has the following requirements:

- 1 Gbps Ethernet physical layer
- Industry-standard SFP module-based fiber interface
- Individual frame jitter that is below 25 ns
- Link asymmetry that is below 100 ns
- Low latency
- Constant link delay (no protected path switching)
- No third-party traffic

Travel time jitter and link asymmetry requirements are the most difficult specifications to meet but are essential for the UHS relays to maintain common time and synchronize individual samples down to the nanosecond level across hundreds of kilometers.

Synchronization could, in theory, be accomplished using two sources of time, supplied independently to the individual relays. Although simple, the independent clock approach does not satisfy a key requirement that UHS relay designers want to achieve, namely, that the relay protection functions be insensitive to external sources of time and any conceivable failure modes associated with those sources. To meet this requirement, relays rely on an internal time source (highly accurate, temperature-compensated crystal oscillator). External sources (when present) are strictly monitored, and if acceptable, are used to discipline the phase of the 1 MHz sampling clock.

When TW87 is enabled, a dedicated communications link allows the two UHS relays to form a strong synchronization bond. A ping-pong message exchange is used to continuously measure the fiber-optic link delay between the two relays. To maintain synchronization, it is essential that the link between the relays provide a constant transport delay.

While not a primary requirement, the TW87 scheme channel latency needs to be in line with the UHS relay operating speed because the excess channel delay directly affects the operating time of the TW87 element.

Using the outlined point-to-point link requirements, a list of communications technologies that can be used to establish the link is shown in Table 1. Using a pair of single-mode fibers is the simplest solution. A pair of fibers is required because normally one fiber is used for data transmission and the other fiber is used for data reception.

Table 1 Point-to-Point Communications Link Options

Transport Technology	Acceptability	Comment
Single-mode fiber pair	Yes	Difficult to justify for sole use by protection
CWDM	Yes	
DWDM	Yes	
OTN	Yes	
SONET / SDH	No	Jitter
Ethernet network	No	Jitter
MPLS	No	Jitter
Software-defined networking-based (SDN-based) Ethernet	Sometimes	Affected by device construction and network setup

The use of wavelength division multiplexing (WDM) is much more common in the industry. In WDM, tightly controlled wavelengths of light (colors) are used to transport multiple communications links over the same fiber. Simpler systems use coarse WDM (CWDM) with up to 16 channels available in the 1,310 nm and 1,550 nm bands. More advanced systems use dense WDM (DWDM) with 44 or 88 discrete channels allocated in the 1,550 nm band. This band has the lowest attenuation and leverages the availability of Erbium-doped fiber amplifiers (EDFAs), thus supporting long-distance transmission. Applications using any of the three intended technologies are shown in Figure 2.

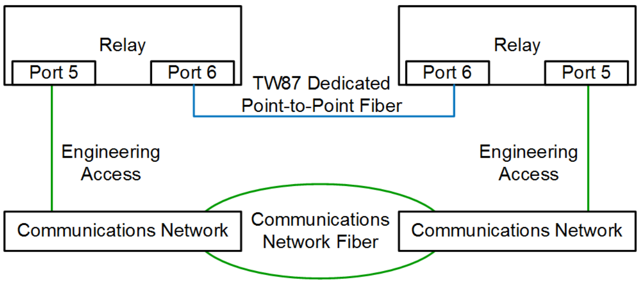

Figure 2. Intended relay application (two systems shown side by side).

OTN is the latest generation of communications technologies, intended primarily for metropolitan aggregation and long-distance core networks. OTN provides a transparent digital wrapper intended for the end-to-end transport of lower-speed network technologies, such as SONET, Carrier Ethernet, MPLS, storage area networks, and so on, and it is currently a preferred access method for large data centers.

SONET, SDH, Ethernet, and MPLS cannot currently be used to meet the TW87 link requirements. SDN can in some cases be configured to provide sufficiently low jitter, but this performance remains manufacturer-specific and design-dependent.

3. Optical Communications Network at PG&E

PG&E is an investor-owned electric utility company serving over 5.4 million electric and 4.3 million gas customers. Headquartered in San Francisco, California, PG&E serves a territory of over 70,000 square miles in northern and central California. In addition to its extensive electric and gas networks, PG&E also operates a large optical network linking virtually all major high-voltage substations in its territory. PG&E’s network uses high-quality single-mode fibers typically deployed along the high-voltage transmission line right of way using optical ground wire (OPGW) technology.

PG&E's communications network uses state-of-the-art OTN technology [3], with 61 nodes already in service. Ten additional nodes are in the works at the northern edge of the service territory, promising full network coverage in the near future. PG&E's utility neighbor to the north, Bonneville Power Administration (BPA), is in the process of commissioning a similar OTN system. Similarly, PG&E's neighbor to the south, Southern California Edison, operates an extensive fiber-optic network with a long history of offering fiber-optic services on the open market.

4. OTN Technology Primer

OTN is a tightly coupled set of technologies aimed at providing transparent, multiservice transport for a wide variety of packet and TDM-based technologies. OTN scales well beyond 100 Gbps per transported tributary and includes a framework for efficient wavelength division control, reconfigurable optical add-drop multiplexing (ROADM), optical amplification, centralized network management, and all optical long-haul transport capabilities. OTN is highly configurable and works hand in hand with the latest generation of SDN.

As an optical transport technology, OTN should be seen as a large number of centrally managed point-to-point links that can be deployed at will. DWDM and the advanced modulation schemes in use today allow the OTN technology to approach the presently known limits of the capacity of single-mode fiber, as shown in Table 2.

Table 2 Single-Mode Optical Fiber Capacity as a Function of Distance

Capacity (Tbps)	Modulation	Typical Distance (km)
4 to 7	PM-BPSK	5,000
8 to 15	PM-QPSK	3,000
12 to 21	PM-8QAM	1,200
16 to 27	PM-16QAM	500

The optical fiber capacities provided in Table 2 assume the use of DWDM with 88 or more individual wavelengths. Each wavelength is modulated using binary phase shift keying (BPSK), quadrature phase shift keying (QPSK), or quadrature amplitude modulation (QAM) with a constellation of 8, 16, or more symbols (8QAM, 16QAM). For modulation rates above 40 Gbps per wavelength, it is also customary to use polarization multiplexing (PM).

Use of multiple wavelengths provides guaranteed (physical) separation of traffic between different streams, while the high transmission speeds make it possible to transparently encapsulate legacy rates and services (such as MPLS or SONET/SDH). OTN transport has standardized the rates shown in Table 3 (applied per DWDM wavelength).

Table 3 Standardized OTN Rates

Signal	Approximate Data Rate
ODU0	1.244 Gbps
ODU1	2.498 Gbps
ODU2	10.037 Gbps
ODU3	40.319 Gbps
ODU4	104.79 Gbps
ODUflex	Any configured rate

Rates are carefully selected to allow easy mapping of various legacy channels, for example, Gigabit Ethernet or OC-48 (2.48 Gbps) SONET services. Slightly larger data rates make it possible to transparently transfer channel timing while at the same time supporting frequency tolerance range and the clock jitter mask requirements.

Although supporting lower rates (such as ODU0), OTN gets in the zone at 10 Gbps or above. For example, at those rates, the ODU2 10.037 Gbps channel can be used to multiplex together one Gigabit Ethernet LAN tributary, one OC-48 (2.48 Gbps) SONET channel, and 5 MPLS channels supplied using 1 Gbps Ethernet ports. All of these tributaries are transported “one bit at a time” (in a round robin fashion, sending one, two, or more bits from each channel, depending on the rate of the tributary). Actual rate management is more complex, with multiple options available to the network engineer, but it is important to note that fine granularity (bit by bit) of such multiplexing allows the subtended systems to maintain their time and frequency synchronization requirements.

OTN systems use forward error correction (FEC), which results in slightly higher rates by the time fully multiplexed electrical signals are delivered to their corresponding modulators and/or WDM channels. FEC provides an additional layer of robustness, allowing the optical transport specialist to monitor the bit error rates (BERs) for each wavelength and proactively adjust the system parameters before the errors become visible to the end customers.

The summary above barely scratches the surface of OTN network capabilities. For additional details, interested readers are directed to [3] and the easily accessible information on the Internet.

5. San Ramon Laboratory Test

To verify the PG&E network’s ability to transport the UHS relays’ TW87 messages, a live equipment test was conducted at the PG&E San Ramon Communications Systems Test Laboratory. Initial conversations among the engineers involved in the project indicated a high level of confidence on the side of the PG&E communications team, caution on the side of the PG&E protection team, and major reservations on the side of the relay design team.

The teams agreed to conduct two independent tests. The first test would use DWDM-capable SFPs plugged directly into the relays. The selected wavelength would then be brought directly into the PG&E OTN platform optical plane using the “alien wavelength” interface module. This approach guaranteed that the UHS relay would have direct access to the optical fiber for the given wavelength. The OTN system hardware used for the test is shown in Figure 3.

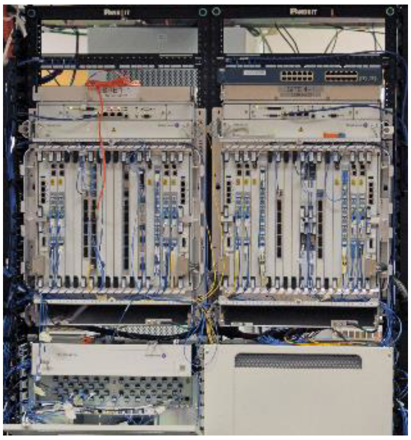

Figure 3. OTN system used during the test.

An alien wavelength card is normally used when the traffic from a neighboring carrier needs to traverse the network without interference. Allocating an entire alien wavelength to a single relay channel in effect means committing a resource with a 100+ Gbps theoretical capacity. In the PG&E network case, resource allocation is slightly less critical because the individual wavelengths are typically modulated at 10 Gbps. The alien wavelength approach still provides much greater efficiency than using a dedicated dark fiber pair (15+ Terabit resource).

A detailed diagram showing the alien wavelength test setup is shown in Figure 4.

A second test involved the use of ODU2 multiplexing. In this test, a 1 Gbps link used by the UHS relays becomes a simple tributary to the ODU2 (10 Gbps) stream, with additional services (including PG&E SONET and MPLS networks) being multiplexed in at the same time. The OTN multiplexing approach provides the best bandwidth utilization, with a 1 Gbps link being provisioned to consume a 1 Gbps resource.

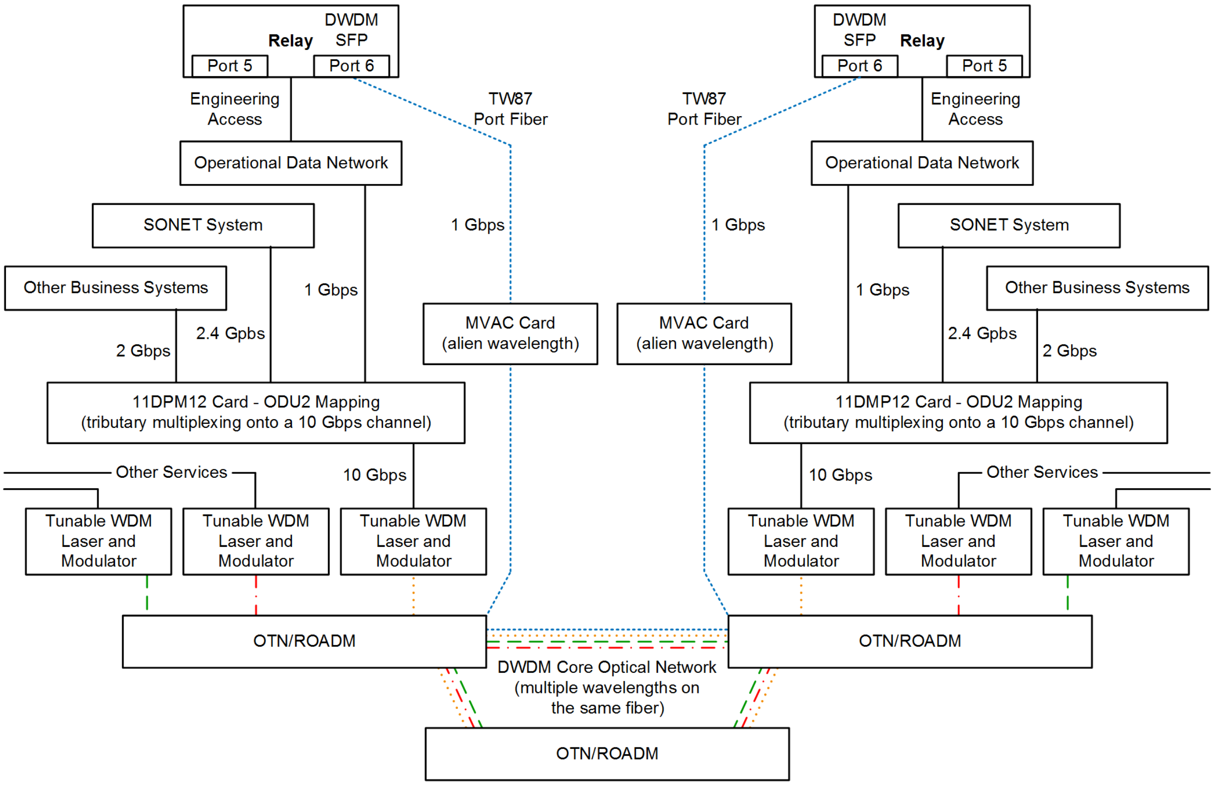

Figure 4. Alien wavelength test setup.

A detailed diagram showing the ODU2 test setup is provided in Figure 5.

The results for the two tests are shown in Table 4.

Table 4 OTN Test Results

Test	Error Correction	Measured Latency (µs)	Measured Jitter (ns)
Alien wavelength	None	1.6	<8
ODU2 tributary	EFEC	60.8	<8

As the results show, the UHS relay link was successfully established in both cases. The alien wavelength test provided the measurement of the actual fiber length connecting the two OTN systems in the lab. The jitter measurement shows that there was no detectable jitter (< 8 ns). The ODU2 multiplexing test shows additional latency caused by the FEC system. This additional latency was expected and is fully controlled by the OTN network engineer.

Available settings and the associated end-to-end delays are shown in Table 5.

Table 5 FEC and Associated Transport Latency

Line-Side FEC	Tributary Rate (GHz)	Line Rate (GHz)	Latency (µs)
None	1.25	10.709	17
RS FEC	1.25	10.709	23.5
EFEC	1.25	10.709	59
EFEC2	1.25	10.709	157

The exact meaning of the various options (RS FES, EFEC, and so on) is beyond the scope of this paper, but can be found in [3]. What matters is the close agreement between the expected and the measured results of the test.

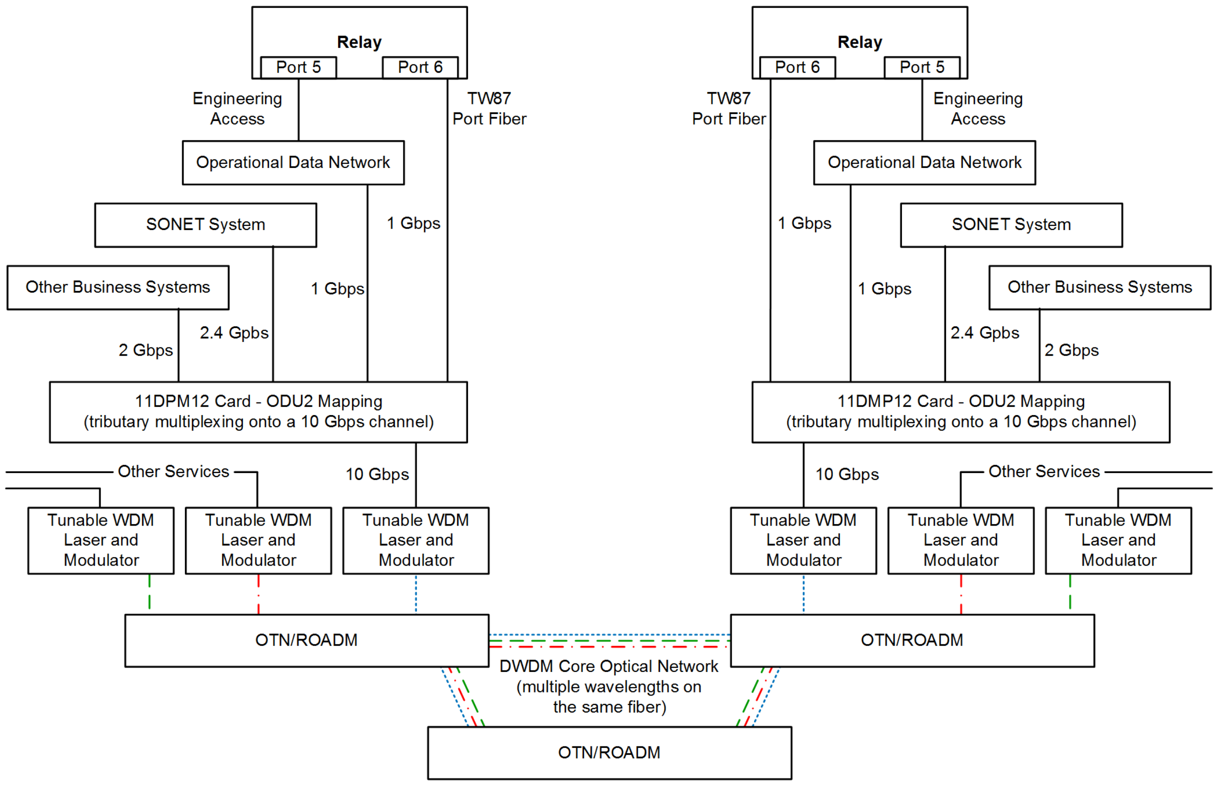

Figure 5. ODU2 multiplexing test setup.

6. Relay Design Opportunities

As explained in the introduction and clearly demonstrated during the test, modern utility communications systems are more than capable of satisfying the protection system data communications needs.

Faced with the continuously evolving communications capabilities, protection device designers are challenged to use the latest Ethernet protocols while at the same time having to support the legacy 64 kbps synchronous interfaces.

Ethernet communications are especially convenient in multiterminal applications in which a single relay needs to exchange data with multiple peers. In order to support mission-critical protection applications, an Ethernet network must be purposefully engineered and capable of providing guaranteed quality of service. General-purpose Ethernet networks are inadequate for this purpose. Instead, fully configured networks must be used, with key technologies including SDN, Ethernet pipe transport over SONET, and time-sensitive networking (TSN).

Packet-based networks are typically associated with queuing delays and an inability to control transport delay variations. As long as the total latency is low enough to meet the protection system requirements, these problems can be solved by adding a network-based time-synchronization service, such as Precision Time Protocol (PTP, also known as IEEE 1588) or one of its profiles (IEC 61850-9-3 or IEEE C37.238).

Ethernet-based protection schemes work as long as the Ethernet network can guarantee that the same time-synchronization signal is delivered to all devices that can communicate with each other, meaning time synchronization and communications should always work in tandem. To meet this goal, time synchronization must become a guaranteed core network service. Traceability of the network time source to Coordinated Universal Time (UTC) is less important, provided the same time is being distributed to all communicating devices.

The exact method for distributing time throughout the network core is less important. It can be proprietary or standards-based as long as the time distribution is reliable and unconditionally cybersecure. One such system based on SONET transport is described in [4]. SONET and SDH are especially convenient for this use because the system cannot operate without establishing a strict synchronous connection with the neighboring nodes.

The system described in [4] is a great example of the time distribution service that demonstrates the level of performance that needs to be provided by any network core. The exact implementation is less important but applies equally well to the OTN network core and the edge networks emanating from that core (e.g., Ethernet, MPLS).

When considering OTN and maximum capacity limits for a single-mode optical fiber, it is important to note that power system protection and control applications do not generate enough traffic to fully load or justify the OTN system. Present day protection and control needs are easily met with a single OTN wavelength. Remaining OTN capacity can be leased or used for other purposes.

7. Conclusion

While not every electric power utility can be expected to switch to OTN technology in the near future, the PG&E and BPA examples can be seen as forerunners of things to come over the next decade. The authors are very excited about the potential of OTN technology and future optical technology advancements.

Relay design engineers are finally finding themselves in an environment with virtually no bandwidth constraints, making it possible to exchange the optimal amount of data required for a given protection application. This means much greater penetration of differential protection schemes, with virtually all key resources protected using some type of optical communications.

Acknowledgements

The authors would like to extend special gratitude to the PG&E protection engineering team, Davis Erwin and Monica Anderson, who are instrumental in successfully evaluating the latest protection technologies at PG&E.

Bibliography

[1] K. Fodero, C. Huntley, and P. Robertson, "Deterministic Communications for Protection Applications Over Packet-Based Wide-Area Networks," proceedings of the 4th Annual PAC World Americas Conference, Raleigh, NC, August 2017.
[2] J. Sykes, D. Day, K. Fennelly, V. Skendzic, and N. Fischer, "Sharing Direct Fiber Channels Between Protection and Enterprise Applications Using Wavelength Division Multiplexing," proceedings of the 44th Annual Western Protective Relay Conference, Spokane, WA, October 2017.
[3] Alcatel-Lucent 1830 Photonic Service Switch, Product Information and Planning Guide, Issue 4, February 2015.
[4] K. Fodero, C. Huntley, and D. Whitehead, "Secure, Wide-Area Time Synchronization," proceedings of the 12th Annual Western Power Delivery Automation Conference, Spokane, WA, April 2010.

20190130 • TP6846