

Automation at Protection Speeds: IEC 61850 GOOSE Messaging as a Reliable, High-Speed Alternative to Serial Communications

Nicholas C. Seeley
Schweitzer Engineering Laboratories, Inc.

Presented at the
10th Annual Western Power Delivery Automation Conference
Spokane, Washington
April 8–10, 2008

Automation at Protection Speeds: IEC 61850 GOOSE Messaging as a Reliable, High-Speed Alternative to Serial Communications

Nicholas C. Seeley, *Schweitzer Engineering Laboratories, Inc.*

Abstract—With the emergence of Ethernet-based, protection-speed protocols, the adoption of such technologies within projects starts slowly and begins to increase as these technologies become more universally accepted. IEC 61850 is the first standard that includes such a protocol, GOOSE messaging, which will become more and more commonplace within the industry as time and the number of projects in which it is used increases.

The advantages of protection protocols transmitted via Ethernet are numerous. A recent project involving intersubstation load shedding over 12 km for a large oil refinery clearly illustrated this. With high-speed information needed for load-shedding functions and system metering value updates acceptable at lower speeds, serial communications dictate the need for separate channels, or the addition of serial multiplexer devices to bridge all types of data, while Ethernet communications allow for various data protocols across the same line more efficiently. The original specification for the project called for serial-based communications using multiplexer technology. However, the details of the project began to show the unfeasibility of this architecture; Ethernet with IEC 61850 GOOSE messaging proved to eliminate all of the problems that serial communications brought about. This paper uses this recent project as the basis for exploring the complexity that serial communications can introduce into a system and how Ethernet, specifically IEC 61850 GOOSE messaging, can be used to simplify the system architecture, yet maintain the speed, reliability, and security at a comparable cost.

I. INTRODUCTION

As technology expands and the speed of communication becomes faster and more reliable, industries that are not in a position to be on the experimental edge begin to adopt the proven. The proven technologies tend to stay static for years before their once experimental phase matures into the robust and reliable phase. This point tends to become a major paradigm shift for companies that have relied on one set of technologies and begin to realize a new set of technologies can accomplish the same job more efficiently. The recent introduction of the IEC 61850 standard has become the latest turning point at which the tried and true serial communications are often successfully supplanted by Ethernet-based technologies. While IEC 61850 has not rendered serial communications worthless, it is certainly advancing the prospect of a global shift within the electric utility community towards adopting Ethernet communications throughout the substation for data acquisition, automation, and some protection functions.

This ideology could be seen first hand during the development of a recent project at a major international oil

refinery. As will be discussed, the preferred serial communications of the initial design became a hindrance, and ultimately unfeasible, considering the restraints imposed. Employing an Ethernet-based technology, namely the GOOSE protocol within the IEC 61850 protocol suite, allowed for complete integration of all required system data, given the physical limitations of the communications architecture.

The project involved implementing a fast load-shedding scheme within the refinery to assure system stability given any event that would cause a loss of generation to the system. In load-shedding schemes, time plays an integral role in deciding if the system will maintain stability or collapse; therefore, the quicker that load can be taken offline to compensate for the lost generation, the more likely the system frequency will recover and ensure that the effects of the event that caused the loss of generation are mitigated.

This particular project involved an interesting array of challenges, mainly because it involved integrating a newly constructed pumping station with an existing station 12 km away. Integrating the new with the old always requires interesting engineering solutions, and in true form, the challenge was how to communicate the needed information from Substation 2 (old substation) to Substation 4 (new substation) in the most efficient manner, while being economically sensitive. The information could be separated into three distinct categories: high-speed data, low-speed data, and engineering access traffic.

A. High-Speed Data

High-speed data involve all information concerned with the trip signals or breaker operations that isolate generation from the system, as well as the trip signals initiated to trip load offline. Consider the operation of a generation breaker. When this breaker operates, generation is immediately lost and the system capacity is reduced by roughly the amount the generator was supplying to the system.¹ Once this generation is lost, it is of the utmost importance that enough load be removed from the system to maintain system frequency. Given this requirement, it becomes obvious that both the indication that generation has been lost and the corresponding trip signals to the loads selected to shed must be transmitted quickly and securely. If the loss of generation is not detected quickly enough or the trip signals are slow to arrive at the

¹ The specifics of load shedding and the calculation of how much load to shed versus how much load is lost are beyond the scope of this paper.

chosen loads, then the system may not be able to recover, or, more likely, an underfrequency backup scheme will operate.²

B. Low-Speed Data

Low-speed data encompass all information that, while important for calculation of various set points within the load-shedding system, does not change frequently enough nor suddenly enough to warrant the need for high-speed transmission. These values include MW flow, disconnect switch status, load consumption, etc. These values are used to arm the load-shedding system, but the high-speed data actually trigger the system to operate. Because these low-speed data play no role in the triggering of the load shedding, they can be updated less periodically.

C. Engineering Access Traffic

Engineering access traffic is information that is not associated with the actual real-time operation of the load-shedding system but permits the retrieval of historical monitoring and configuration information from the system. These data are not speed-critical but can require a higher bandwidth due to the amount of information being transferred. Frequently used engineering access traffic includes event report retrieval, Telnet access to individual relays within the system, remote desktop services, and ad hoc diagnostics.

D. Combining the Three Types

Traditional serial communications would require three separate communications channels to transport the three separate types of necessary data. Not to mention, when dealing with systems that are as critical as a load-shedding system, redundancy is nearly always required. This means a minimum of six communications channels with the proper redundancy must be available for the load-shedding system to be functional. For communications channels within a substation, adding more channels is as easy as running a cable to each device. When building a substation from the ground up, proper planning and system specifications make this issue a triviality compared to the overall system. However, as was the case with this project, communications between substations become more of a challenge.

With two substations separated by 12 km, more communications channels mean more cable runs or fiber over those 12 km. This is often not feasible, and, as was the case with this particular project, only two pairs of fiber were allotted for all communications between the two substations. Given this restriction and the original intent of using high-speed serial communications for the high-speed data, the obvious choice was to install a multiplexer to combine the different data and send one data stream across a pair of fiber. Using this method, installing two multiplexers at each end carrying identical data streams across the two pairs of fiber allowed for redundancy of the system. See Fig. 1 for the serial

communications architecture of a nonredundant system. A redundant system would have a total of two multiplexers on each end, and each device would be connected to both.

II. SERIAL COMMUNICATIONS

While this architecture was, at least on paper, a viable option, an uncomfortable uncertainty existed with regards to the use of multiplexers. Two concerns immediately came to the surface—determinism and reliability. Reliability always plays an important role when it comes to designing systems, especially systems for refineries, where massive amounts of money are dependant, to a large extent, on the ability to keep the electricity flowing. Multiplexers are largely used for communications purposes that can rarely be classified as critical and are designed and manufactured accordingly. To that extent, serial multiplexers are efficacious in what they are designed for, but what they are designed for, in most instances, does not include high reliability within extreme environments. Not to mention, with the current trend of Ethernet and other high-speed data communications, serial communications products are becoming less and less available.




Fig. 1. Nonredundant System Architecture

Determinism also became a very important consideration. It is extremely important that when high-speed data are sent from one end or the other, the data arrive at the other end 100% of the time, without worry of retransmissions, being held in queue, or simply getting lost or corrupted. If high-speed information is sent, then it must, under all circumstances, make it to the other end reliably and in a timely fashion. This takes extensive testing and collaboration with the multiplexer manufacturer in order to guarantee the performance required for the system.

Considering all of these requirements, multiplexers have still found their way into the utility protection arena and have performed as well as could be expected. However, Ethernet communications have begun to reach a maturity where utilities are more accepting and able to rely on their performance, and, therefore, are making them a larger piece of their system's communications schemes.

² Underfrequency backup schemes are very common and recommended in load-shedding schemes for the added assurance that if the load-shedding system is inoperable, a backup system is in place. Coordination of the load-shedding system with the underfrequency scheme is also of great importance, but outside the scope of this paper.

III. ETHERNET COMMUNICATIONS

While serial communications remain widely used throughout the communications world, Ethernet communications are becoming more prevalent for substation communications. Serial communications are sure to remain because dedicated point-to-point, high-speed, secure, low-overhead protocols are still the preferred standard for protection communications. However, Ethernet is taking hold in this realm as well.

The IEC 61850 standard includes a high-speed, multicast protocol: GOOSE messaging. GOOSE messaging is a non-routable, OSI Layer 2, broadcast/subscription Ethernet-based protocol that evolved from the UCA 2.0 GOOSE messaging protocol. GOOSE, while not as secure as some of the more common point-to-point protocols, is still very useful in some protection-type applications. In particular, it is ideally suited for this load-shedding application, as it is sufficiently fast and, being an Ethernet protocol, can run on the same communications line with several other protocols.




Fig. 2. Ethernet Network

IV. THE DECISION BETWEEN SERIAL AND ETHERNET

The proposed load-shedding system called for redundant communications. For this redundant communications ring, we were provided with two pairs of optical fiber. With this in mind, we were restricted in that all communications had to use one communications path. Considering this, we were faced with two options, multiplexing the serial data or switching to Ethernet.

It is important to note that other options do exist, namely wireless, for the transmission of data over 12 km. The two substations are within line-of-sight with no obstructions, so this would be an ideal application for spread-spectrum radio or similar technologies. However, in this particular case, the area may encounter sand storms, which would be more than enough to disrupt the signal and stop communications. Considering this possibility, the system architecture was limited to wired communications solutions only.

Economically speaking, the multiplexed serial data would have required the purchase of four multiplexers, as opposed to Ethernet, where there were already switches installed at each substation. However, each switch would need to be supplied with cards for 9-micron fiber-optic cable connections. Regardless, the Ethernet cards were roughly one-fourth of the cost of the multiplexers. In terms of reliability, the addition of more hardware inherently decreases the reliability of the system.

While the addition of the 9-micron Ethernet cards to the switch was technically no different than adding multiplexers to the system, in terms of the hardware added reducing the reliability of the system, the Ethernet cards are rated for extreme environmental conditions, whereas the multiplexers are not. In the end, in terms of both reliability and cost, moving to Ethernet communications not only became a feasible alternative, it looked to be a better alternative than the original design, considering the limitation of the physical communications paths that were allotted.

V. SYSTEM ARCHITECTURE

The system is segregated into two halves, local and remote. As mentioned earlier, the remote substation is located 12 km away from the local substation. The load-shedding system algorithm is centralized on a computer with a Linux[®] operating system, referred to hereafter as the LSP (load-shed processor), at the local substation. Data collected from the field intelligent electronic devices (IEDs) are concentrated in a communications processor and sent via unsolicited messaging to the LSP. These data consist of the low-speed data discussed earlier, breaker and disconnect switch statuses, and meter analog values. These data are gathered by the LSP and used to perform system calculations to decide if generation is lost on the system, how much, if any, load should be shed, and which loads are selected. Low-speed data (data sent by the communications processors) are essentially used to calculate the reaction in the event of lost generation. High-speed data communicate what event has occurred (the tripping of a generation breaker, tieline, etc.) and send the commands to trip the required load.

Because the LSP resides in the local substation, all relays communicating these high-speed “event” data can communicate serially. Pre-made fiber-optic patch cable can be used between the local relays and the LSP, making it possible to connect all the relays providing high-speed data. The need for the high-speed Ethernet GOOSE protocol became evident when gathering and transmitting high-speed data, along with low-speed data from the remote substation. Because these low- and high-speed data, along with Telnet-type engineering access traffic, can exist on the same communications line, Ethernet is the prime choice for this application. See Fig. 3 for the Ethernet system architecture.

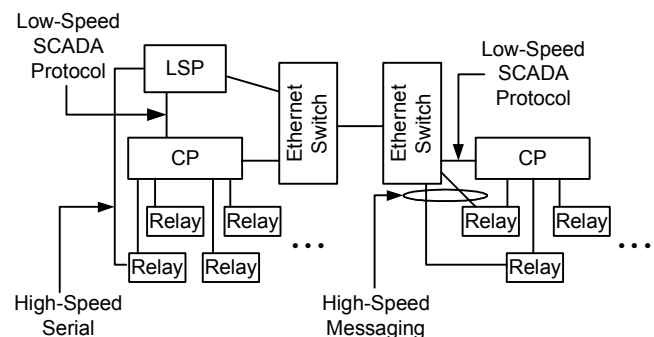


Fig. 3. Ethernet System Architecture

The sequence of events for a typical load-shedding event would initiate upon the opening of a generation breaker or the receipt of a trip signal from the tripping relay associated with a generation breaker. This breaker status, or trip status, would be sent high speed, serially in the case of an event occurring in the local substation and via Ethernet GOOSE in the case of the remote substation, and then received at the LSP. The LSP receives and processes this signal and issues **TRIP** commands to the relay outputs of the loads that have been selected to shed. Fig. 4 is typical of the local substation where the high-speed serial communications are used.




Fig. 4. Basic Function of the LSP

Fig. 5 illustrates the path the trip signals originating from the remote substation must follow. Because of the intermediary Ethernet link, the data path is not as direct as within the local substation. This Ethernet segment, while still fast enough for our application, does slow the overall response of the load-shedding system. This issue will be addressed later.




Fig. 5. General Architecture

A. Why Timing Is Critical

Load shedding is becoming a more popular protection-type functionality. Until relatively recently, load shedding was not capable of being done quickly by today's standards, but faster computer processors and data communications have allowed industries to begin investigating subcycle load shedding.

The purpose of load shedding is to protect a system in the event of lost generation and help the system to maintain stability and system frequency. When generation is lost, if the remaining generation is not capable of outputting additional power to make up for the deficit, the system frequency will eventually decay beyond recovery and will collapse. However, if a system is in place that can calculate exactly how much

power will be lost and how much the remaining generation can supply in excess of what it is currently supplying, it can then calculate how much load must be shed in order for system frequency to maintain stability. This is the essence of a load-shedding system: to calculate the effect a loss of generation would have on a system and then determine how much, if any, load should be shed to maintain stability. Therefore, load shedding is acting in a protection role, more so than a traditional automation role. This concept begins to blur the lines of automation and protection. In a modern, high-speed load-shedding system, the load-shedding algorithm is being processed quickly enough that dynamic system decisions can be processed and operate in times under 16 ms. Not only can it be processed in under 16 ms, but the trigger that initiates the process, the process itself, and the receipt of the output of the process are transmitted in less than one power cycle. Traditional automated systems would require significantly longer times to process data and make a dynamic decision based on these data. Technology now allows us to not only process this information but communicate it to the necessary hardware, which could be separated by larger distances, to take corrective action. Such technological capability becomes a veritable panacea for power stability related issues.

In the case of load shedding in particular, the load-shedding system must be coordinated with backup under-frequency protection schemes so that the system frequency never falls below the underfrequency threshold. The time that it takes for the system frequency to decay to this point is largely dependant on the system makeup. Systems with large synchronous machines will have the benefit of their system frequency being held up by the sheer inertia behind those machines. Therefore, the frequency will not decay as quickly as a system with smaller generators and motor loads, as well as resistive loads, which results in a slower frequency decay.

One of the interesting concepts to note in this discussion, however, is that the line between automation and protection is being blurred, whereas, not very long ago, automation and protection were completely separate functions. An automation group at an industrial plant or utility would work on applications involving remote control of the system and data acquisition, and the protection group would focus on high-speed protection of the system assets. Now automation schemes are being coordinated with protection schemes, bringing both groups together. In the case of load shedding in particular, the protection engineers complete detailed studies of how robust the system is in response to a loss of generation. The automation engineers use that information to determine what to expect of each generator when writing the LSP algorithms or vice versa.

B. Time Tests

The load-shedding system presented here has two components: a serial side and an Ethernet side. Below we will compare the performance of the two different communications types and discover how they compete with each other. As mentioned before, the Ethernet communications scheme does

add time to the performance, which should be no surprise because it is an additional path that the serial communications side does not encounter.

For an illustration of the test setup, refer to Fig. 5. In Table I, input IN102 detects a trip signal from a generator breaker. IN102 is part of a GOOSE messaging data set that triggers on change. The changing state of IN102 causes a change of state GOOSE message to be issued. The receiving device has mapped the GOOSE message to local bit RB02, which is subsequently mapped to the high-speed serial protocol transmit bit. This bit is transmitted to the LSP where it is then processed and, in turn, issues a **TRIP** command via the high-speed serial protocol. That **TRIP** command is received by the sending device and mapped to a GOOSE message to be transmitted over the Ethernet network. Once transmitted, the receiving device detects the GOOSE message, processes it, and asserts an output to trip the selected load. This whole process takes, as shown in Table I, an average of 40 ms, roughly two-and-a-half cycles. From input powered to 80% nominal to output contact 90%, conducting the entire process takes approximately 40 ms. Observed timing with time-synchronized Sequential Events Recorder (SER) records have ranged from 35 ms to 42 ms.

TABLE I
ETHERNET PATH
(REFER TO FIG. 5. TIMING INDICATIVE OF AVERAGE TIMES RECORDED)

Action	Time Duration Since Previous Action	Time Duration Since Start
Trigger and GOOSE Message Publication at GOSP2 AC1	Start	Start
Wide-Area GOOSE Trigger Transmission		
GOOSE Trigger Message Receipt at GOSP4 AC	11 ms	11 ms
GOOSE-to-Serial LSP Interface		
Subsequent Serial Message Publication to LSP Within GOSP4 AC	4 ms	15 ms
LSP Algorithm Processing		
Receipt of Serial Message From LSP at GOSP4 AC	12 ms	27 ms
Wide-Area GOOSE Trip Transmission		
GOOSE Trip Message Receipt at GOSP4 AC	10 ms	37 ms
Trip Control Output at GOSP2 AC2	4 ms	41 ms

In Table II, we cut out the Ethernet side of the communications and rely completely on the serial communications. See Fig. 4 for a basic illustration of the test setup. An input is received and transmitted via a high-speed serial protocol to the LSP. The LSP processes the input and issues a **TRIP** command. The **TRIP** command is received by the tripping device and asserts an output. Taking out the Ethernet loop, we see a greatly improved performance. We measure 13 ms from

input to LSP decision to output. With the direct serial communications, we were well under a cycle.

TABLE II
SERIAL PATH
(REFER TO FIG. 4. TIMING INDICATIVE OF AVERAGE TIMES RECORDED)

Action	Time Duration Since Previous Action	Time Duration Since Start
Trigger and Serial Message Publication at GOSP2 Relay	Start	Start
Wide-Area Serial Trigger Transmission		
Serial Trigger Message Receipt at GOSP4 LSP	5 ms	5 ms
LSP Algorithm Processing Plus Wide-Area Serial Trip Transmission		
Receipt of Serial Trip Message From LSP at GOSP2 Relay AC	4 ms	9 ms
Trip Control Output at GOSP2 Relay	4 ms	13 ms

VI. CONCLUSIONS

Ethernet communications are becoming a viable option for protection-related automation schemes. While, with this scheme, the direct serial communications operated in one-third the time the Ethernet scheme took to operate, it should be noted that at this time there is no direct interface for GOOSE messaging into the LSP system. Because of this, the GOOSE message had to trigger a high-speed serial message that could be processed by the LSP and then sent back via high-speed serial, which triggers a GOOSE message to be received by the relay shedding the load. Taking this inefficiency out of the equation takes 8 ms out of the transfer time for the Ethernet-based scheme. This brings the time comparisons a little closer, where the serial communications system operates in less than one cycle, the Ethernet system operates in approximately two cycles. Both are excellent, and the Ethernet system adds flexibility to the system and requires fewer communications lines because it is part of the Ethernet network. However, the obvious drawbacks are slower responses and less security.

While it is important to address the options available within certain constraints and how this project is similar to any number of projects currently under development, it is interesting to note the paradigm shift that is occurring within the industry. Technology is amalgamating the automation and protection roles into one superset; the line is being blurred, and the results will be coming in shortly.

VII. ADDITIONAL READING

R. Jenkins, D. Dolezilek, and M. Gugerty, "Case Study Comparison of Serial and Ethernet Digital Communications Technologies for Transfer of Relay Quantities," proceedings of the 33rd Annual Western Protective Relay Conference, Spokane, WA, October 2006.
Available: <http://www.selinc.com/techprsr.htm>.

VIII. BIOGRAPHY

Nicholas C. Seeley graduated from the University of Akron in 2002 with a B.S. degree in Electrical Engineering. After graduation, Nic began working at American Electric Power in Columbus, Ohio, for the Station Projects Engineering group where he focused on substation design work. In June 2004, Nic was hired at Schweitzer Engineering Laboratories, Inc. in the Systems and Services Division as an automation engineer where he has been involved in the development, design, implementation, and commissioning of numerous automation-based projects.