Station Device and Network Communications Performance During System Stress Conditions

Karen Leggett, Roy Moxley, and David Dolezilek
Schweitzer Engineering Laboratories, Inc.

Presented at the
Southern African Power System Protection Conference
Johannesburg, South Africa
November 10–12, 2010

Previous revised edition released July 2010

Originally presented at the
1st Annual Protection, Automation and Control World Conference, June 2010

Station Device and Network Communications Performance During System Stress Conditions

Karen Leggett, Roy Moxley, and David Dolezilek, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Automation and protection systems that go beyond the most basic operations require some form of communication between relays and a central processor and between relays themselves. Typical practices in evaluating or specifying the communications capabilities of relays or networks measure the time for a signal to transport from the internal processes of one device, through the network, to the output of another device. While this may be an important measure, it may not reflect operating conditions in the power system.**

This paper describes research to measure the complete system time to transmit, communicate, receive, and process digital messages in a system used for any automated control or protection purpose. In addition to measuring single bits of data, tests were performed using multiple bits and "data storms" that could be seen during significant system events or disruptions. Complete tests included measurement of processing times within relays, as well as network times to distribute data to multiple devices through point networks.

As Ethernet communications increase the ability to exchange information between numerous devices in a station, care needs to be taken to evaluate the system as a whole and not just as a group of points. This paper provides methods to evaluate overall system performance during realistic conditions that could be encountered within installations during test and once in service. Understanding the impact of message size, transmission rate, and multiple device transmissions on system performance are all crucial to optimize system design and set realistic performance expectations.

I. Introduction

The use of discrete on/off signals in protection systems has traditionally been accomplished through contact inputs and outputs wired into a protection scheme. Digital relays made it possible to use advanced communications systems to provide virtual direct inputs to a scheme from a remote location, (telecommunications) without the need to install discrete field I/O wiring or deal with the delays associated with conveying status information through physical means. As microprocessor-based relays integrated multiple functions into one physical device, many communications protocols were developed to potentially communicate virtually thousands of pieces of information from each intelligent electronic device (IED). These protocols include independent standards, such as IEC 60870 and DNP3, managed by a committee (users group) funded by a collection of vendors and users that organize enhancements and testing. The standards also include National Institute of Standards and Technology-approved protocol standardization created by a standards-related organization (SRO) and offered via a "reasonable and nondiscriminatory" license, including proprietary protocols [1], as well as other standards such as IEEE C37.94. In general, the proprietary protocols used in industrial applications do not satisfy power system requirements for interoperability and resiliency. Interoperable solutions have evolved into serial channel networks and Ethernet networks.

II. Virtual Field Wiring Cables Via Digital Communications

Serial SRO protocols travel over serial communications cables and act as a virtual field wiring cable to communicate discrete and analog field measurements via multiple virtual wires within the cable. These serial SRO protocol cables are dedicated to this peer-to-peer communications channel and do not share bandwidth with other communications [2]. This fixed bandwidth provisioning guarantees that these serial SRO protocols will not be affected by additional network traffic or network segment failures. The logical data flow path is the same as the physical serial cable path, peer-to-peer directly between two IEDs connected by the serial cable.

Two major differences between messaging based on Ethernet networks versus direct serial channels impact design, performance, diagnostics, and upgrade of communications systems. Unlike direct serial channels, the physical path of Ethernet network messages does not match the logical path. Ethernet messages that logically pass directly between two peers actually physically pass through several cables and Ethernet switches. Further, the physical path will change over time without the peers' knowledge. Also, Ethernet protocols are either routed to a network address (not a specific device address) or multicast throughout the entire network so that data created by one IED can be sent to many. Multicast means that the message has no destination address, cannot be routed, and must be sent to every port and device on the Ethernet network, even to network segments where it is unwanted but unstoppable.

Ethernet protocols like DNP3 LAN/WAN (local-area network/wide-area network), Modbus® TCP (Transmission Control Protocol), Telnet, FTP (File Transfer Protocol), IEEE C37.118 for synchrophasors, IEC 61850 MMS (Manufacturing Message Specification), Generic Object-Oriented Substation Event (GOOSE), and Sampled Values all share the available bandwidth on an IED or switch port. Switches attempt to overcome the very real possibility of message latency via several "best effort" quality of service technologies [1]. One of the multicast protocols within IEC 61850, GOOSE, is often used for Ethernet virtual cabling. It is designed with a fixed publication rate if no data are changing, immediate publication upon change of state, and a

brief period of rapid publication after the change of state. These behaviors increase the likelihood that subscribers will quickly learn of data change, even if the network drops messages and then slows down the publication rate to reduce network load and switch and IED processing.

Ethernet network connections allow several different virtual channels and multiple protocols to all share bandwidth on the IED Ethernet interface. However, commissioning, verification of performance, and troubleshooting are made difficult, because it is not evident from the physical Ethernet cable connected to the IED what logical point-to-point connections and multicast publications and subscriptions are active. Perhaps equally important is the uncertainty for a technician who wishes to take a relay out of service but no longer has physical terminations, test switches, or wiring diagrams. The technician wants to prevent unwanted interruption of virtual cables but has no idea where they are being used in the network. It is also impossible to disable only one of several logical connections within one physical Ethernet connection.

Finally, adding IEDs to an existing network requires thorough knowledge of existing communications behavior. Configuration changes made to an IED IEC 61850 GOOSE may very well change the contents and therefore the virtual wiring terminations among subscribing IEDs. Much the same as untested changes to physical field wiring terminations at an IED, each new GOOSE message configuration may change the virtual interconnections.

III. Multiple Protocol, Numerous Application, Shared Bandwidth Ethernet Networks

Communications systems for digital information using a direct serial connection between relays provide for fixed latency and a monitored communications channel. The need for consistent, high-speed communication, preferably with a fixed latency and a monitored channel, is inherent in providing a high-quality protection system. While direct peer-to-peer communications provide these two features, as soon as a network, particularly an Ethernet network, is involved, there are complications. Not only are IEC 61850 GOOSE messages widely used for telecommunication of field wiring values, they also share protection information between multiple devices in a substation via digital communication (teleprotection). IEC 61850 protection signals, other messages via additional IEC 61850 protocols, and other non-IEC 61850 protocols all share the bandwidth of the Ethernet network and support different applications.

IEC 61850 classifies application types based on how fast the messages are required to be transmitted among networked IEDs [3]. The standard also specifies the performance of each type of application, documented as time duration of message transmission. Table I lists the message types.

The IEC 61850 communications standard documents different performance classes for different message types within the suite of IEC 61850 protocols as shown in Table I [3]. IED networks are expected to maintain the level of performance documented in Table I constantly, regardless of network application and messaging activity.

TABLE I
IEC 61850 Message Types and Performances [3]

Type	Applications	Performance Class	Requirements (Transmission Time)
1A	Fast Messages (Trip)	P1	10 ms
		P2/P3	3 ms
1B	Fast Messages (Other)	P1	100 ms
		P2/P3	20 ms
2	Medium Speed		100 ms
3	Low Speed		500 ms
4	Raw Data	P1	10 ms
		P2/P3	3 ms
5	File Transfer		≥1000 ms
6	Time Synchronization		(Accuracy)

IV. Application Transmission Times

The time duration to create and deliver messages between IEDs via a protocol is the message transmission time, represented in Fig. 1 by $t_{transmission}$ (transmission time) = $t_a + t_b + t_c$. The time duration to publish information in IED1 (Physical Device 1), deliver it via a protocol message through the network, and act on it in IED2 (Physical Device 2) is the information transfer time represented by $t_{transfer}$ (transfer time) = $t_{transmission} + t_{f2}$. This information transfer time duration is the time truly useful to the design engineer, because it represents actually performing an action as part of a telecommunication automation or teleprotection scheme. Transfer time, $t_{transfer}$, is easily measured as the time difference between the accurately time-stamped Sequential Events Recorder (SER) records in IEDs with synchronized clocks.

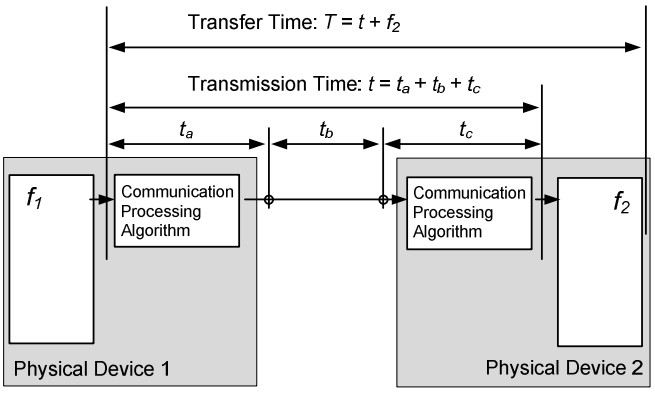

Fig. 1. Application transmission time definition (from IEC 61850-5)

V. Ethernet Network Latencies

LANs used to interconnect substation IEDs, servers, concentrators, and gateways are built with Ethernet cables and substation-rated, managed Ethernet switches. Ethernet uses shared bandwidth provisioning techniques to merge all the

message packets of multiple conversations onto various network segments [1]. The network devices use variable provisioning and path routing techniques, which increase the likelihood that packets will safely navigate the network. However, these same techniques make the network activity uncertain and nondeterministic, which is generally why the network is reflected by a cloud. The message enters the cloud and eventually exits at the destination, but it is not clear or consistent how the message will navigate the network each time. All of these network elements and media introduce latencies that need to be analyzed in order to match the substation automation system requirements.

Within flexible bandwidth provisioning networks, there is a network saturation point (e.g., 80 percent of bandwidth usage) where this competition for bandwidth noticeably slows down message transit through the network. However, if IED network topologies are chosen carefully and are sufficiently small, message latency through the network cables and switches will not be statistically significant to the teleprotection application transmission time of 3 milliseconds [1]. From Fig. 1, ($t_{transmission} = t_a + t_b + t_c$) time through the network, t_b, is considered sufficiently small and unchangeable for wire-speed switches. Also, once the network is built, protection designers have no influence over the performance of network latencies. Therefore, logic processing and network processing in the IED will have the greatest influence on $t_{transmission} = t_a + t_b + t_c$.

VI. Ethernet Network Compensation Methods

Because multicast messages are nonroutable and unstoppable, IEC 61850 adopted Ethernet network compensation methods to improve message delivery [1]. The unique nature of Ethernet dictates that only one connection and one data flow path be active at a time for any IED, regardless of how many physically redundant IED connections and network paths exist. Instead of redundancy, IEEE 802.1 relies on failure and recovery techniques to reestablish broken links. In an effort to use off-the-shelf Ethernet products, IEC 61850 continues to use general-purpose Ethernet switching technology and recommend that network switching technology and redundancy be added into the IEDs to compensate for the failure/recovery and multicast features. This has the unintended consequence of burdening the IEDs with network compensation processing, in addition to their IED network processing and primary function, such as protective relaying.

As previously mentioned, multicast behavior means that each time a multicast message, such as GOOSE, is received on a switch port, it is automatically sent to every other port. The unneeded but unstoppable messages increase the burden on the switch, waste bandwidth, and increase latency of the necessary GOOSE exchange. In addition, the IED processor burden increases because the IED must process both the unnecessary and the necessary GOOSE messages.

Each time an IED receives a multicast or broadcast message, it has to decode the message and see if it should process it. The IED examines the multicast address, data set reference, application ID, and GOOSE control reference of each message to verify that it is the correct message from the correct IED. If it matches the IED subscription configuration, the IED processes the contents and maps them to internal memory. If it does not match, the message is discarded after the verification processing.

One of the techniques to alleviate the network burden of multicast or broadcast messages is the virtual local-area network (VLAN).

IEEE extended the Ethernet Standard 802.1 with the designator Q for message quality, which includes extensions for optional VLANs via a previously unused field in the Ethernet header tag. IEEE 802.1Q VLAN, or QVLAN, can divide a physically connected network into several VLANs, as shown in Fig. 2. QVLANs originated from a need to segregate network traffic from different departments inside one enterprise. While keeping sensitive information private, QVLAN techniques can be used to restrict traffic flow of multicast or broadcast messages to a single QVLAN and therefore the devices within it.

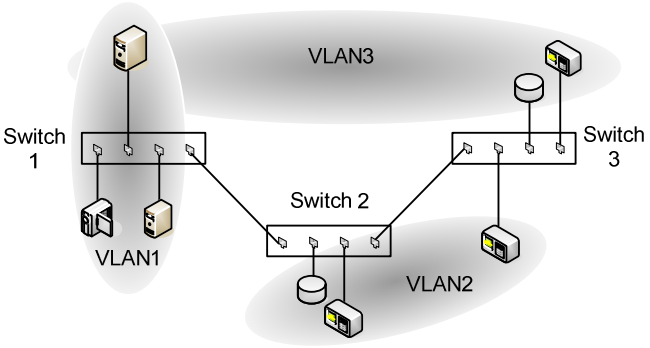

Fig. 2. Switched Ethernet and QVLAN configuration

IEC 61850 adopted the use of QVLAN tags to identify multicast messages and overcome the inability to perform network routing by performing manual routing. Because of the unwanted and unstoppable automatic distribution of multicast messages, the manual routing is performed in reverse. The multicast messages are routed everywhere but are only allowed to pass through ports from which they have not been blocked. In IEC 61850 networks, QVLAN tags are implemented within the multicast message by the publishing IED, potentially used by switches for manual routing, and are ignored by the subscribing IEDs. This is one of several network processing tasks that have been forced into the IEDs to compensate for inadequate data flow capabilities in Ethernet networks. Switches unable to perform QVLAN filtering or those configured incorrectly will not work correctly and may block even wanted GOOSE transfer. Best engineering practice methods within IEC 61850 dictate a unique QVLAN identifier, an IEEE 802.1Q tag, for each GOOSE message publication. Each GOOSE message becomes a virtual cable whose contents are virtually wired to each IED that needs the data, but no others. This is done by creating a VLAN between the source IED and the destination IED(s).

GOOSE has become an efficient method of using digitized communications to replace the traditional field wiring technique of physical copper conductors conveying state or analog information between a sensor and IED. A GOOSE message acts like a virtual cable, with information from several conductor pairs (virtual wires) within it. The QVLAN becomes the unique cable designator. Ethernet switches use the QVLAN to cause the Ethernet network to act as power system engineers wish and guide the GOOSE virtual cable to only those IEDs that need it. Network designers add settings to each switch port to identify which QVLANs to allow and which to restrict. Though configuration-intensive, this mechanism helps mitigate the wasted bandwidth, transit delays, and unnecessary IED processor burden associated with unrestricted multicast. Like many aspects of Ethernet, the promiscuous nature of sending all multicast messages everywhere until told to stop is the opposite of what protection and automation engineers want. These engineers would prefer that virtual cables go nowhere until told to do so. Also, when unexpected multicast traffic is added in the future, it will result in wasted bandwidth, transit delays, and unnecessary IED processor burden if it has no QVLAN tag or a QVLAN tag in which the ports were not set to anticipate and restrict. This will happen any time a new device is added intentionally or when an unwanted or unexpected device is added without the designers' knowledge.

Another compensation technique adopted to reduce transit latency of multicast messages due to network congestion is to use priority tagging, per IEEE 802.1p. In order to compensate for the bandwidth-sharing techniques of Ethernet, packet prioritization was created to emulate long-standing SRO serial protocol message prioritization methods. In this case, each packet, regardless of the protocol within it, can be assigned a priority. This is done similar to QVLAN within a previously unused field in the Ethernet header tag. It is another of several network processing tasks that have been forced into the IEDs to compensate for inadequate data flow capabilities in Ethernet networks. For switches and IEDs that support the feature, the priority tag indicates the importance of each packet relative to the others. Packets with highest priority are sent to the top of the queue. If a lower-priority message is in process, or packets with the same or higher priority are in queue, even prioritized packets must wait.

Unlike QVLAN, if a switch does not support priority or is configured incorrectly, it will not prohibit message transit through the network. However, it will not prevent unwanted transit latencies by treating all messages the same during a transmission backlog. Perhaps more importantly, potential message latency due to incorrect use of the priority tag may not be evident during normal operation of the network. Latencies may occur only during times of power system and Ethernet network stress, long after commissioning testing, and at the time when latencies are most dangerous.

The only effective method to segregate Ethernet multicast traffic and GOOSE virtual cables is to follow these simple rules:

- Assign each GOOSE virtual cable a unique QVLAN tag.
- Allow no multicast messages on the network without QVLAN tags.
- Assign each GOOSE virtual cable an IEEE 802.1p priority tag.
- Disable all unused switch ports.
- Configure every switch port to block delivery of every multicast message to the connected IED except the QVLAN virtual wires that the IED has subscribed to within its configuration file.

VII. Network Processing Impact on Teleprotection via IEC 61850 GOOSE

Now let us consider a teleprotection scheme where all GOOSE messages are top priority. The standard gives requirements for transmission time for an application. From Table 1, for Fast Message Performance Class P2/P3 teleprotection, the entire transmission time, including network latency, is required to meet 3 milliseconds. Though IEC 61850 does not mention that it is acceptable to exceed this time for any reason, IEC 60834-1 requires that 99.99 percent of commands for intertripping protection schemes be delivered within 10 milliseconds. Clearly, design for reliability requires that we assume that each multicast is the initiate message for a teleprotection application and needs to meet the documented performance criteria. Further, in a multipurpose protection scheme, we now need to understand how multiple simultaneous teleprotection applications working through the Ethernet network will influence the performance times of each network. Again, though not addressed to date by any standards, design for reliability requires that we also assume that each multicast of each simultaneously active teleprotection application is the initiate message for that teleprotection application and needs to meet the documented performance criteria of 3-millisecond transmission time.

In the station of Fig. 3, consider the protection messages that may be shared between devices. By a protection engineer's definition, these messages are all considered the highest priority:

- External fault detection for bus and line protection blocking.
- Breaker failure transfer tripping [4].
- Recloser control based on line or feeder status.
- Single-pole trip declaration.

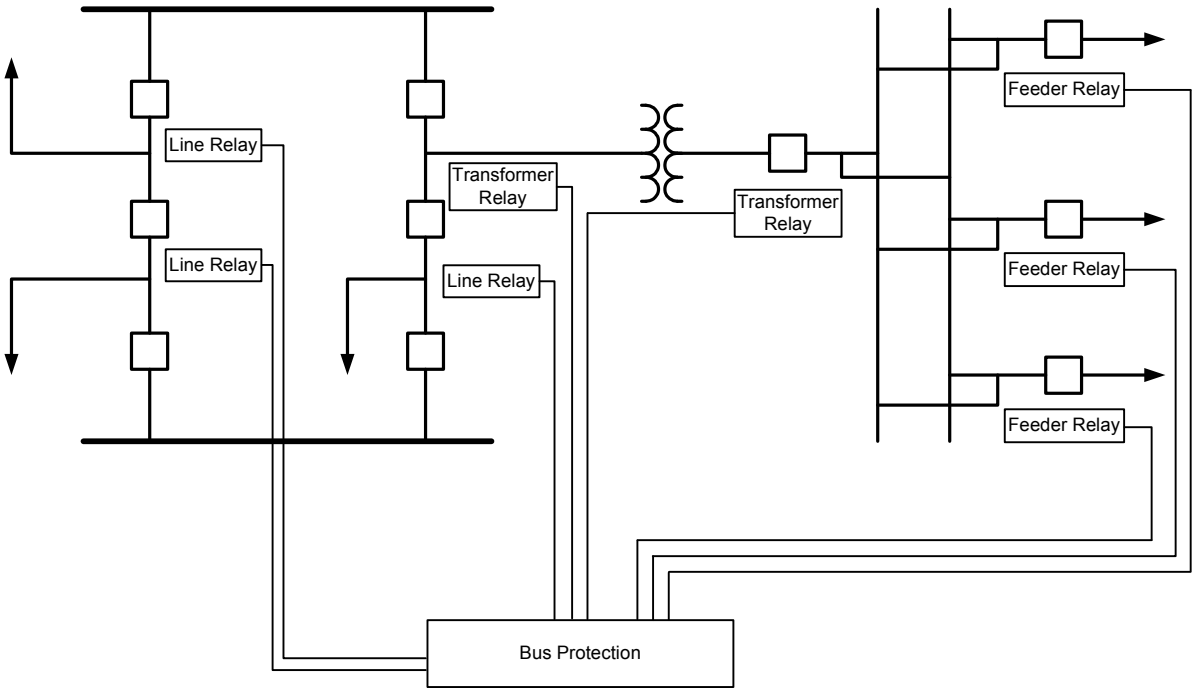

Fig. 3. Station one-line diagram with protection

In addition to the multicast messages conveying these signals, there are permutations that will add to the publication, transportation, and message receipt processing burden. For example, each of the line and feeder relays may be configured to convey directional elements in the same or different messages. Separate messages for internal and external fault detection can be configured. A communications status message can be created by each IED to tell other networked IEDs of detected failure of an incoming GOOSE or test mode status for the different applications [3]. Obviously, adding each of these new messages from each IED increases the use of the available Ethernet bandwidth. Even if all the signals are conveyed via a single GOOSE message with all of these signals in a single payload, it will be published more frequently due to more signals changing state.

Taking these possible messages and applying them to the Fig. 3 system, any fault, either on a bus or line, will produce either forward or reverse fault declarations from most, if not all, of the relays at the station. The bus relay may send trip signals to each of the relays on a bus section. If each relay sends only two messages, all the relays will need to receive and process 18 messages. Three messages per relay, certainly within the realm of possibility even without breaker failure, would lead to 27 simultaneous messages. Applying the QVLAN to each GOOSE message in the IED and correctly configuring the Ethernet switch provide the ability to filter unneeded GOOSE messages at the switch port. Even so, and even if the system is carefully designed such that power system disturbances do not result in massive simultaneous multiple message transmission, the system designer is left with no means to guarantee that the natural retransmit interval for GOOSE messages will not result in many messages being transmitted just before a fault. Such a condition results in the relays receiving many "background" messages just before they receive important tripping or blocking messages.

VIII. Testing Message Transmission Time

Protection designers are interested in satisfying the transfer time for multiple concurrent GOOSE teleprotection schemes. It is likely that telecommunication for automation via GOOSE is happening concurrently as well. Therefore, a test was created, as illustrated in Fig. 4, to measure the transfer time for teleprotection applications during various network message-loading scenarios. Seven different relays from three different manufacturers were used for the tests. Multiple multicast scenarios represent simple message payload with only change of state information, payload with change of state plus analog values, and various frequencies of data change to simulate different event activities.

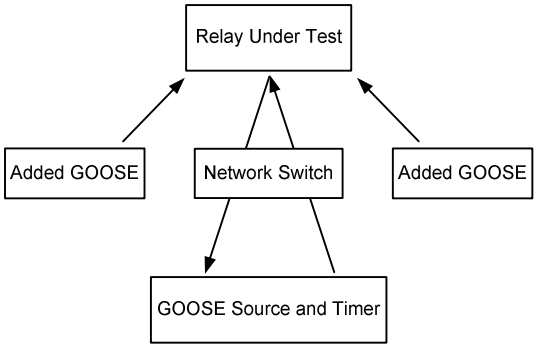

Fig. 4. GOOSE performance test setup

It is important to note that not all IED manufacturers support the IEEE 802.1Q and IEEE 802.1p compensation techniques of QVLAN and priority tags within the published GOOSE message. Therefore, the performance test was staged

with all messages published on the same QVLAN and with the same priority. This not only served to support the same subscription and publication configuration within IEDs from each vendor but also to demonstrate the impact of not configuring QVLAN filtering in the Ethernet switches. In addition, a GOOSE message is published immediately after one of the values in the payload changes state or passes through a dead band. Repetition of the multicast continues at a maximum rate after the change, and then if the payload remains unchanged, the repetition rate slows until it reaches a preconfigured minimum repetition. The multicast will continue to publish at this slower rate until another data change occurs. For this test, data in the added GOOSE messages are changed at the same rate as the maximum multicast repetition, and so the degradation of the repetition rate for these GOOSE does not occur and does not influence the times measured in this test.

The relay under test is at the top of Fig. 4. A specially configured timer relay, shown at the bottom of the figure, was used and was the same in each test. The ultra-high-speed timer relay publishes a GOOSE test start (TS) message with an 8-bit payload. The timer relay starts an internal timer and changes all 8 bits simultaneously, which results in an immediate publication, followed by the repetitious messages. The relay under test subscribes to this TS multicast, maps the 8 received bits into internal logic, and performs protection logic to map these same 8 bits to outgoing bits. The relay under test publishes these 8 bits in a relay under test GOOSE message. The timer relay subscribes to this relay under test multicast and stops the internal timer when the internal logic receives the bits, indicating the roundtrip exchange for the 8 bits.

The network switch is a modern, 100 Mbps, wire-speed device. As previously discussed, the switch speed is not ever expected to contribute significantly to any of the times measured during testing.

Teleprotection schemes rarely require the exchange of 8 bits either direction, and the simultaneous change of state of these 8 bits represents nontypical primary equipment activity and network traffic. However, this test loading is considered the base, or quiescent, test case in order to investigate performance of a teleprotection application during a period of existing network activity. The first test used no additional messages or network traffic. The results of ten separate performances of this quiescent test case are shown in Fig. 5.

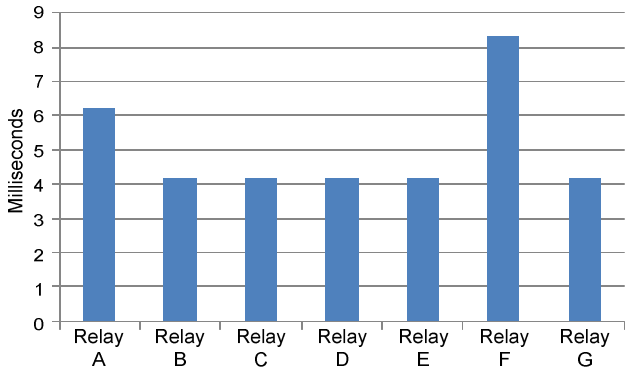

Fig. 5. Average quiescent test case results

Note that the times shown represent one complete message roundtrip: $(t_{transfer1} + t_{transfer2}) = (t_{transmissionTR} + t_{f2} + t_{transmissionDUT} + t_{f1})$. Recall from Table 1 that the maximum allowable transmission time is 3 milliseconds for trip messages, so the allowable roundtrip processing time from the results of Fig. 5 is two one-way times of 3 milliseconds, plus two 2-millisecond functional processing intervals, for a total of 10 milliseconds. If the average quiescent roundtrip time of less than or equal to 10 milliseconds was the sole design criterion for a teleprotection application via GOOSE, then all of the relays would pass the test. However, even with no additional network traffic, it should be understood that Ethernet communications and, more importantly, some IED Ethernet processing are not deterministic. By automating and repeating the Fig. 4 test, the roundtrip maximum time was measured for thousands of tests. While typical (and even average) times were very fast, it is interesting that the maximum times were significantly higher than average in some cases, as shown in Fig. 6. Because GOOSE works as a subscription to a repetitious multicast publication, the delays may be caused by slow processing of messages or the loss of messages. This test measured the roundtrip time of status indication transfer and did not monitor message loss, nor was it determined what percentage of the test cases went beyond the 10-millisecond maximum roundtrip time.

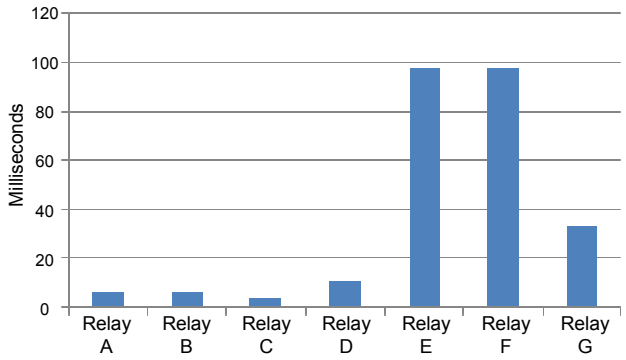

Fig. 6. Worst-case quiescent test case results

These worst-case times may not represent a problem for some virtual wiring applications. However, they demonstrate nondeterminism and may impact critical protection operations, wide-area protection and remedial action scheme (RAS) systems, and system stability. It is important to remember that this test illustrates that received message processing can have a large impact on the total transfer time.

As discussed previously, there are many cases where relays will see and be expected to process messages from multiple applications simultaneously. To test for these conditions, increased data transfer was simulated by connecting another GOOSE message source (described in Fig. 4 as Added GOOSE) that was configured to publish a GOOSE AD1 message with 16 Boolean bits that changed every 2 milliseconds. The relay under test was configured to subscribe to the AD1 multicast and map the 16 bits to internal resources. The roundtrip measurements were repeated. Fig. 7 represents the worst-case roundtrip time from at least 2,000 repetitions for each relay under test. Again, because of

the nature of multicast, the delays may be caused by slow processing of messages or the loss of messages. This test measured the roundtrip time of status indication transfer and did not monitor message loss, nor was it determined what percentage of the test cases experienced such long maximum roundtrip excursions.

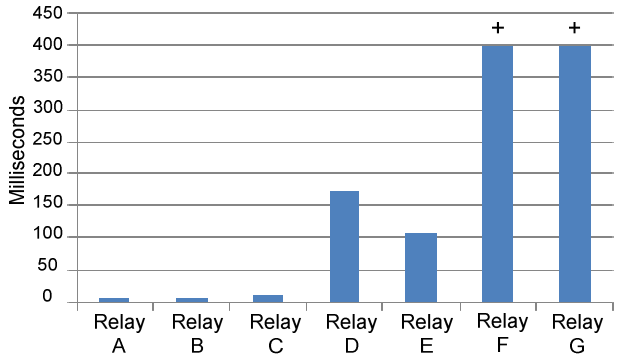

Fig. 7. Additional 16 subscribed bits, worst case

In this test, over half of the relays tested had at least one worst-case roundtrip time that far exceeded the maximum documented in Table 1. Two of the relays had at least one roundtrip time longer than 0.4 seconds, which is truncated in Fig. 7 to be equal to 0.4 seconds. If the relays under test do not need to subscribe to AD1 for designed applications, or if those values typically change much less frequently, this test represents a simulated data storm or incorrect network configuration. Both the unnecessary message processing and unnecessary payload content mapping to internal logic simulate additional network processing stress in the IED. Incorrect configuration should be detected and corrected at commissioning, and a data storm probably will not last longer than a few tenths of a second.

It is important to note that the timer relay performance was verified to exhibit deterministic behavior empirically as it was Test Relay B in every test case. It was documented to behave as consistently as both the timer relay and the relay under test in Fig. 4. Also, for some relays, the roundtrip times are unchanged from the typical quiescent test, further indicating that neither the timer relay $t_{transmissionTR}$ nor t_{fl} contributed to longer roundtrip measurements.

The second Added GOOSE device is introduced, injecting an additional GOOSE AD2 message with another 16 bits, changing every 2 milliseconds, for a total of 32 additional subscribed bits in each relay under test. At this point, it can be seen in Fig. 8 that all of the relays, except A and B, resulted in at least one roundtrip time longer than four cycles.

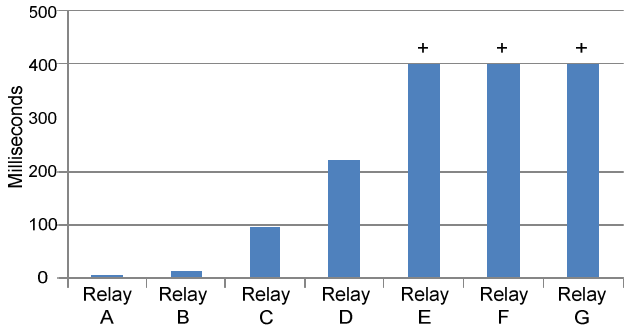

Fig. 8. Additional 32 subscribed bits, worst case

At this scale, it is difficult to tell, but the Relay A roundtrip time has not changed from the typical quiescent. And, once again, because of the nature of multicast, the delays may be caused by slow processing of messages or the loss of messages. This test measured roundtrip time of status indication transfer and did not monitor message loss, nor was it determined what percentage of the test cases experienced such long maximum roundtrip excursions. And, as previously discussed, if the relays under test do not need to subscribe to AD1 and AD2 for designed applications, or if those values typically change much less frequently, this test represents a simulated data storm or incorrect network configuration. Both the unnecessary message processing and unnecessary payload content mapping to internal logic simulate additional network processing stress in the IED.

All of the test results shown in Fig. 5 and later were sorted based on the results shown in Fig. 8. Note that the results increase from left to right on this test case. Notice also that the shape results from the other test cases do not correlate well with the final test case. This demonstrates that performance cannot be predicted during a data storm by testing in the quiescent case.

All of the GOOSE messages sent in the tests described above were subscribed to by the relay under test. The next test simulated an incorrectly configured switch that allowed unsubscribed multicast traffic to pass through to the IED port. Four large GOOSE messages with complex payloads were added to the network via an additional relay, not shown in Fig. 4. No relays subscribed to the four new large payload GOOSE messages, nor did the relay that published them subscribe to any messages. Results are shown in Fig. 9 for at least 2,000 interactions on each relay under test.

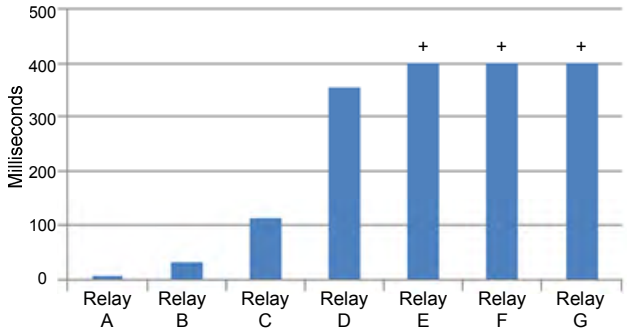

Fig. 9. Additional 32 subscribed bits, plus four unsubscribed messages

Note that when a switch is configured incorrectly and unsubscribed messages reach the relay, the relay performance, in most cases, degrades in response to the additional processing. Roughly 80 percent of the network processing in the relay must be performed before the relay is sure that the message does not meet the subscription criteria and should be discarded. Relay A performs excellently and still has the same worst-case roundtrip times as reported for the quiescent condition in Fig. 5. In the case of Relay A, a patent-pending system was employed within the relay to segregate and filter GOOSE traffic inside the relay, as shown in Fig. 10.

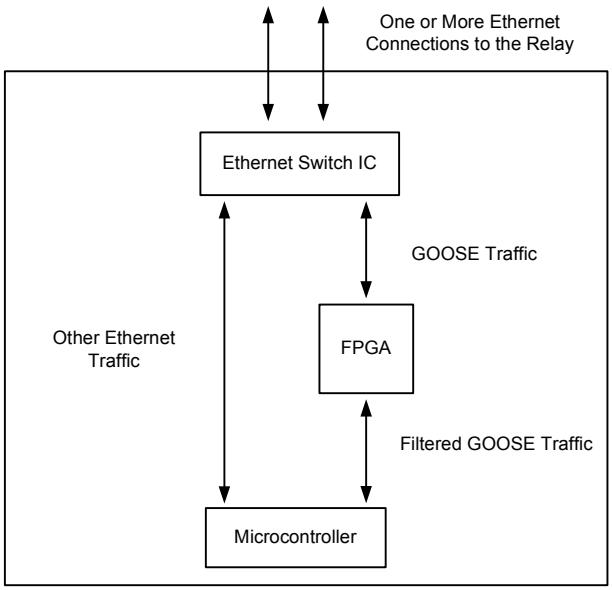

Fig. 10. GOOSE message segregation prior to processing

In this device, GOOSE traffic is segregated from other lower-priority traffic and is also filtered prior to any processing. In this way, no processing time is used on unsubscribed messages, improving the response speed for important messages.

IX. Network Mitigation

The first line of defense against the type of data storms that can have a major impact on scheme performance is to design the network to avoid them. In this case, the network includes all the connected relays exchanging GOOSE messages. Recognizing the limitations of all connected relays is the first step. If the relays cannot discriminate between subscribed and unsubscribed messages prior to processing the message header information, then the total aggregate of messages that the relay may possibly receive must be less than what a critical relay can process without delay. Experience has shown that using configuration software capable of importing and displaying IEC 61850 configuration for all the devices, from each manufacturer, on the network will help ensure proper coordination of each device, as well as avoid network problems. It is essential to be able to use such an engineering tool to view the payload contents, publication parameters, and QVLAN tag of each GOOSE message in order to predict network messaging behavior [5].

X. Understanding In-Service Relay Performance

The relays tested had GOOSE transfer times that met the high-speed requirements of IEC 61850 for the quiescent or unloaded states. The key finding of the testing performed was that when network traffic increased, even traffic of unsubscribed messages, it was possible for transfer times to be notably degraded. Recognizing this is the first step to avoiding the problem. Testing relay performance under high network traffic is necessary to verify that the relays will perform as expected.

XI. Conclusions

Relay systems have traditionally been tested under conditions that represent expected system operation. With traditional, hard-wired communication between relays, a system operation test with current injection to the relays was sufficient to verify coordinating times and scheme performance. The introduction of nondeterministic communication, such as Ethernet, into protection schemes must change how scheme tests are performed. New design and testing procedures must be created to verify scheme performance under the worst-case conditions that can be reasonably (and perhaps unreasonably) expected. The lack of correspondence between quiescent and high-traffic relay transfer times means that care must be taken before designing a system based on published specifications. Relays must be designed to operate and transmit data effectively under all possible network traffic conditions, and networks must be designed to ensure no traffic conditions can exist beyond relay capabilities. Finally, due to the failure/recovery nature of Ethernet and the nonroutable and unstoppable nature of multicasting, network compensation techniques in the IEDs are essential. The only effective method to segregate Ethernet multicast traffic and GOOSE virtual cables is to follow these simple rules:

- Assign each GOOSE virtual cable a unique IEEE 802.1Q QVLAN tag.
- Allow no multicast messages on the network without QVLAN tags.
- Assign each GOOSE virtual cable an IEEE 802.1p priority tag.
- Disable all unused switch ports.
- Configure every switch port to block delivery of every multicast message to the connected IED except the QVLAN virtual wires that the IED has subscribed to within its configuration file.

Finally, it is clear that lack of determinism is still a very real concern for use of Ethernet networks for real-time, mission-critical telecommunication and teleprotection applications. When designing a network, recognize that IED processing changes dramatically based on the subscribed and unsubscribed multicast messaging received on its Ethernet port. Due diligence and design for reliability dictate that designers investigate IED processing capabilities during quiescent and high network load scenarios. They need to identify typical transmission time, worst-case transmission

time, the quantity of test cases that exceed the maximum transfer time and by what margin in order to understand what percentage of multicast exchanges do not satisfy the application performance requirement.

XII. References

[1] D. Dolezilek, "Using Information From Relays to Improve the Power System – Revisited," proceedings of the Protection, Automation and Control World Conference, Dublin, Ireland, June 2010.

[2] E. O. Schweitzer, III and G. W. Scheer, "System of Communicating Output Function Status Indications Between Two or More Power System Protective Relays," U.S. Patent 5,793,750, August 11, 1998.

[3] D. Hou and D. Dolezilek, "IEC 61850 – What It Can and Cannot Offer to Traditional Protection Schemes," proceedings of the 35th Annual Western Protective Relay Conference, Spokane, WA, October 2008.

[4] E. Atienza and R. Moxley, "Improving Breaker Failure Clearing Times," proceedings of the 36th Annual Western Protective Relay Conference, Spokane, WA, October 2009.

[5] V. Flores, D. Espinosa, J. Alzate, and D. Dolezilek, "Case Study: Design and Implementation of IEC 61850 From Multiple Vendors at CFE La Venta II," proceedings of the 9th Annual Western Power Delivery Automation Conference, Spokane, WA, April 2007.

XIII. Biographies

Karen Leggett received her BSCSE from Oregon Institute of Technology. She has experience in automation and integration in the nuclear industry and in the pulp and paper industry, and she spent eight years as the SCADA and automation engineer at an electric utility. She is currently the lead integration and automation engineer for distribution products at Schweitzer Engineering Laboratories, Inc.

Roy Moxley has a BSEE from the University of Colorado. He joined Schweitzer Engineering Laboratories, Inc. (SEL) in 2000 and serves as marketing manager for protection products. Prior to joining SEL, he was with General Electric Company as a relay application engineer, transmission and distribution (T&D) field application engineer, and T&D account manager. He is a registered professional engineer in the state of Pennsylvania and has authored numerous technical papers presented at U.S. and international relay and automation conferences. He also has a patent for using time error differential measurement to determine system conditions.

David Dolezilek received his BSEE from Montana State University and is the technology director of Schweitzer Engineering Laboratories, Inc. He has experience in electric power protection, integration, automation, communications, control, SCADA, and EMS. He has authored numerous technical papers and continues to research innovative technology affecting the electric power industry. David is a patented inventor and participates in numerous working groups and technical committees. He is a member of the IEEE, the IEEE Reliability Society, CIGRE working groups, and two International Electrotechnical Commission (IEC) technical committees tasked with global standardization and security of communications networks and systems in substations.

© 2010 by Schweitzer Engineering Laboratories, Inc.
All rights reserved.
20100729 • TP6436-01