Network Errors and Their Influence on Current Differential Relaying

Solveig Ward
Quanta Technology

William Higinbotham
RFL Electronics, Inc.

Abstract

While a point-to-point fiber is the preferred communication link for a protective relaying engineer it is not always available. An OPGW (Optical Ground Wire) on the power line itself could provide a dedicated fiber pair for protection, but often the utility communications network does not follow the power system topology. In addition, multiplexing data over that fiber pair optimizes the use of available communications media. The same physical fiber pair that carries a single channel 64 kbps relaying data could be used for up to OC-192 (10 Gbps) or GigaBit Ethernet.

To use an available communications network for relaying makes sense from an economical standpoint. However, even with digital communications, the network is not error free. While many links are optical fiber which are immune to electrical interference, most networks incorporate wired links or microwave paths in their design. Attenuation on long fibers, synchronous clock inaccuracies causing frame slips, etc., contributes to the errors. For evaluation, protective relays are subjected to extensive bench or simulator testing with injected currents and voltages it is less common to include the communications link in the test. Typically, the relays have a direct connection of the communications ports back to back. This excludes the vital communications link that is integral to the operation of a current differential relay. The results are thus valid for a direct fiber application but not necessarily so over networked communications.

While a relay engineer typically is not directly involved in the design of the utility communications network, the ability to explain how and why critical pilot relaying channel requirements may differ from other data communication needs could go a long way towards obtaining reliable communications services for protective relaying.

This paper examines the type of errors that can be expected over a modern digital communications network and what the consequences for current differential relaying may be. It looks at factors influencing protection system availability and how problems could be mitigated by relay design, proper network configuration and circuit conditioning. Actual testing with current differential relays and a BER (bit error rate) generator illustrate the points.

Power Systems Communications

Protective relaying is one of many services provided over the utility communications network.

The requirement of high speed data transfer for relaying is recognized by telecoms specifications. Protective Relay service requires the shortest Process Response Time of all power system communication services and unlike most other communication requirements, requires short, deterministic, symmetrical latencies. It is also apparent that protective relaying is just one of many services provided by the communications network, and most likely one of the smaller services with regards to amount of data transferred. It might therefore be difficult to justify optimization of the relaying channel in the same way as when a dedicated relay channel was used. Still, this should not be difficult to do, if the teleprotection and pilot relaying demands are fully understood.

While relaying is the most time critical and the most system critical service, the bandwidth required is small. It may be hard to

justify specially designed circuits and communications devices for this small traffic burden, especially if the particular relay requirements are not known or well understood.

The first telecommunication channels deployed for relaying were voice channels over analog microwave and analog phone lines. As telecommunications evolved into the digital era, relaying channels followed. There is now pilot relaying performed over digital phone lines (via CSU/DSU's), and over E1/T1 and SDH/SONET networks via multiplexers. The latest addition to the telecommunication industry, Ethernet, is the next challenge for pilot protective relaying.

Current Differential Relay Communications Requirements

Data is intrinsically different from voice. By definition, data is information which originates in the form of digital representation (binary 1s and 0s) and therefore does not need to be converted to digital within the network, unlike voice that originates from the telephone microphone as an analogue signal. Data originates from terminals, or from a laptop or computer, control monitors and protective relays. Applications generating data from these devices include the sending of emails, transfer of information between files and other forms of enquiry, remote control of machinery and exchange of information between protective relays over a pilot channel. There are several key differences between the characteristics of voice and data and hence the different requirement for successful communications. To complicate matters, pilot protective relaying communications has requirements that are a mix of data and voice. By nature, pilot communications is data but has to operate in real-time, as does voice. However, the relatively high error tolerance and relatively moderate latency requirements for voice are not acceptable for pilot relaying.

The ability to intermix voice and data and preserve the transmission characteristics of each is a primary requirement of T1 multiplexers. Voice transmission, for example, can tolerate a surprising number of bit errors and not significantly affect the quality of the voice at the receiving end. Speed of delivery is important as any delay is noticeable in voice conversations. In telephone company channel banks, multiplexers data transmission is treated just the opposite. Error free transmission is a higher priority than speed of delivery. It is this practice that makes most Telco grade T-1 systems unusable for critical real-time applications such as current differential and phase comparison relaying.

Table I. Communications Requirements

	Data	**Voice**	**Pilot Relay**
Delay (latency) tolerance	High	Moderate/Low (50 ms)	Very low (<20 ms)
Asymmetry and variation in delay tolerance	High	Moderate	Very low
Stream/burst transmission	Bursts	Stream	Stream
Error tolerance	Low	High	Very low
Packet/data loss tolerance	Moderate, by the application requesting retransmission	Some data loss is acceptable until voice quality becomes too low	No
Interruption tolerance	Yes, by the application requesting retransmission	Moderate (100 ms)	None/very low
Protocol standard	Proprietary / standardized	Standardized	Proprietary

Dedicated Optic Fiber

A dedicated fiber pair (dark fiber) is the preferred communications link by the relay engineer as it resembles the conventional point-to-point pilot wire or Power Line Carrier link. An optic fiber pair available for exclusive use by the relays provides optimal performance for digital communications. Dedicated fiber gives a fast and error-free point-to-point connection. The main drawback is that a fiber cut will cause channel

interruption for a long period of time, and many utilities lack expertise and equipment for replacing and splicing a damaged fiber cable.

For a communications engineer, however, to use a fiber pair for 64 kbps, or even less, is a waste of capital resources. The same physical fiber can carry multiplexed data up to OC-192 (approximately 10 Gbps) or GigaBit Ethernet.

Time Division Multiplexing (TDM)

T1 is a term for a digital carrier facility used to transmit a digital signal at 1.544 megabits per second. T1 is used for a wide variety of voice and data applications, embedded in the network distribution architecture as a convenient means of reducing cable pair counts by carrying 24 voice channels in one 4-wire circuit. T1 multiplexers are also used to provide access to higher order transport multiplexers such as 'SONET'.

A T1 frame consists of 24 eight-bit words plus a framing bit. Each timeslot of the frame contains 8-bits of binary information. Each timeslot is called a Digital Signal Zero (DS0) which is sampled 8000 times per second. This sampling rate was chosen because it can adequately represent voice characteristics of a human speaker when using Pulse Code Modulation (PCM). Therefore, each DS0 contains 64kbps (8k samples/sec x 8 bits/sample) of user information. Time Division Multiplexing (TDM) is used to combine 24 DS0's into one T1 frame. Since there are 24 DS0's in a T1 frame, the effective data rate is 1.536 megabits per second. Each frame contains one framing bit, which is used primarily for synchronization. This bit adds an additional 8kbps of overhead to the frame thereby primarily for frame synchronization. This bit adds an additional 8kb/s of overhead to the frame, increasing the information rate from 1.536 Mb/s to 1.544 Mb/s. This 1.544 Mb/s is commonly referred to as a Digital Signal One or DS1. Note that the word T1 and DS1 are used interchangeable, however this isn't really accurate. A T1 refers to the digital transmission system, which happens to operate at DS1 rates.

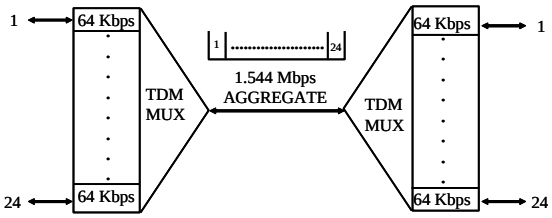

Figure 1. Time Division Multiplexing

TDM (Time Division Multiplexing) networks are synchronous. In synchronous communication characters are spaced by time, not by start and stop bits. There is a constant time between successive bits, characters or events. The timing is achieved by sharing a common clock source.

Transmission is made at a constant bit rate, with a constant frame size. While data transmission is more efficient than asynchronous communications (no overhead start and stop bits are required), a channel fully occupies its time slot whether data is transported or not.

Most protective relaying communications today uses T1/E1 and SONET/SDH technology. This is a synchronous communication using Time Division Multiplexing (TDM) in which the communication link is providing a constant data stream with a fixed bit rate. The higher the bit rate, the higher the bandwidth, and the higher the amount of data that can be transmitted. The data stream is divided into blocks, channels or time slots, each slot being dedicated for a pre-defined function regardless of if the function is transmitting any data or is idle. The major advantage of TDM for relaying is its deterministic latency characteristic.

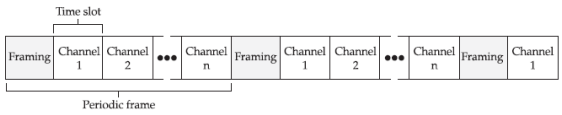

Figure 2. Synchronous communication

Ethernet Communications

Ethernet communication products and applications are becoming more entrenched as the new standard for data transport. Carriers are cutting costs by eliminating the need for overlay networks that require different equipment for different service types. The primary goal for carriers is to accommodate all of these services on a single network and Ethernet IP (Internet Protocol) appears to be the clear winner. The main advantage with this technology is that carriers will no longer incur the added costs associated with conventional TDM networks. In general, Ethernet networks cost less to build due to lower cost equipment and they handle data transfer more efficiently than the traditional circuit-switched, time-division-multiplexed (TDM) alternatives. Ethernet is more efficient because it only uses bandwidth when data is transmitted. TDM systems use full bandwidth whether data is transmitted or not.

Connecting traditional voice, video and data over Ethernet networks has become an attractive alternative to running parallel voice and data networks. It saves money on call and leased-line service charges, while consolidating management, cutting maintenance costs, and increasing user productivity. This is achieved by converging two important traffic types onto one infrastructure and takes advantage of the simplicity and efficiency of IP routing and Ethernet switching.

Ethernet is a packet based technology. The data stream to be transmitted is divided into packets and then each packet is delivered whenever there is bandwidth available. While this results in a more efficient use of bandwidth, there is a certain randomness introduced. End-to-end delays may not be constant. The packet has many alternative routes to reach its destination. The number of nodes and queuing time at each node play a role. Sequential data is typically chopped up into several packets that have different delivery routes. The receiving application has to put them back into correct order.

There are technologies that address these shortcomings and provide solutions for "real time services" such as VoIP (Voice over IP) and Time Division Multiplexing over IP.

Another advantage with Ethernet is the plug-and-play nature of network. A TDM network needs to have carefully designed channel allocation to optimize the bandwidth usage. Ethernet is largely self-configured in this respect.

As modern numerical current differential relays are made for TDM networks, typically using a 64 kbps channel, this paper deals with that technology only. However, network errors discussed here are also valid for a future generation of current differential relays designed for Ethernet network communications.

Communications Considerations

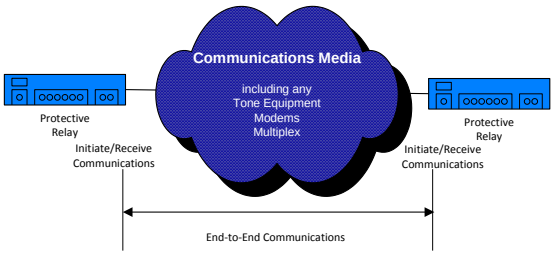

Figure 3. Protective relay communications

A recent report by UTC Research identified reliability as the number one criteria for utility communications networks. Other important factors were high bandwidth, very low latency, ubiquitous coverage, tight security and uninterrupted power supplies.

When a digital communications system is used for teleprotection or pilot protection, the dependability and security of the communications network has to be taken into account for overall protection system reliability.

Of particular concern for relay communications over digital channels is timing issues:

- End-to-end delay; excessive delay due to intermediate devices
- Variable delay; change in delay time from one time period to another
- Asymmetry; different transmit and receive delay paths
- Interruptions and re-synchronization following a switching operation on the network

While a relay may work perfectly over a dedicated fiber, availability of the protection scheme when communicating over a network may not be as high. The reasons are often hard to pinpoint but are likely due to incorrect configuration of intermediate devices (multiplexers) or network errors and switching operations requiring re-synchronization of the protective relay device.

We also have to remember that substation multiplexers reside in a harsh environment. Power system faults cause transients and ground potential rise. Any device that forms a part of the protective relaying system clearly cannot be allowed to be rendered non-operational at the time it is needed.
IEEE 1613 Standard Environmental and Testing Requirements for Communications Networking Devices in Electric Power Substations define two classes.

- The class for devices carrying critical relaying data has the same requirements as protective relays when it comes to SWC, EMI, RFI, etc. The device is not allowed to suffer loss of data during these tests.
- The lesser class allows interruptions but require automatic recovery. In practice this means that the power supplies have to be substation hardened as they otherwise might be damaged by transients or ground potential rise and will not be able to recover.

Timing Issues: Buffering

All data communication devices need to use buffering to ensure dependability. Buffering is used to overcome slight frequency differences and to absorb jitter and wander as these would cause unbuffered systems to overflow or underflow.
Typically, at the minimum, two frames are buffered. This is the case for a substation T1 multiplexer that has been designed and optimized for protective relaying traffic.
Telecom channel banks may use longer buffering, and that can become an issue. While increasing dependability for general purpose data communications, it may defeat a current differential scheme due to excessive delays.
Buffering also contributes to asymmetric end-to-end delay. Two DS-1 frames take 250 us, and varying buffer alignment can cause a difference between transmit and receive delays. Most current differential relays can accept such an asymmetry (0.25 ms) but not all designs. A telecom channel bank with buffering longer than 250 us may cause proportionally higher asymmetry which could disqualify the device for use with many current differential relays.

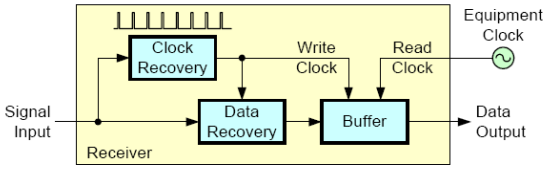

Figure 4. Buffering

SONET networks are made resilient by the use of rings. The most common is the unidirectional ring which means that all traffic in the ring proceeds in the same direction. For an adjacent pair of nodes, if transmit data goes point to point then the receive data must traverse the entire rest of the ring. This will result in significant asymmetric delay. The degree of asymmetry depends on the size of the network; how many nodes that the signal needs to pass though and the distance of the fiber links.
There are technologies that route transmit and receive in the same direction (bi-directional rings) and also methods to ensure that both paths are switched at the same time in case of fiber breaks. However, for a network designed

with unidirectional rings (that work fine for the vast majority of all traffic) it may be hard to justify a specially designed service for relaying only.
The most economic solution is to examine the network characteristics and select protective relaying devices that are suitable for it.

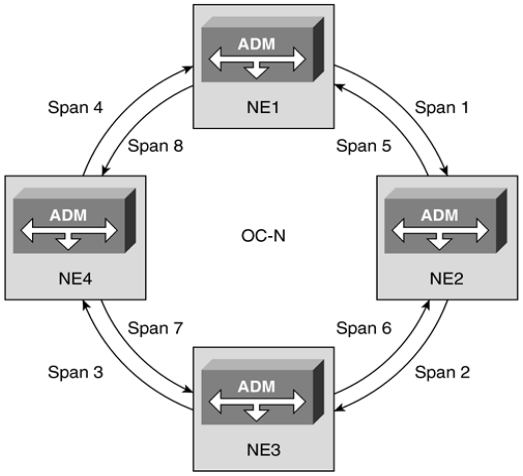

Figure 5. Network ring topology

Network Errors

Digital communications networks are not error free. Errors are caused by transients induced in metallic wires (or devices), equipment failures (a faulty device may send out a garbled bit stream), temperature variations, changing atmospheric conditions of microwave links, etc.
Even the Geostationary clock produces clock errors due to a slight shift of location. Clock errors are also caused by temperature changes where the framing bit wanders with respect to the central clock. Each piece of equipment that a clock transits through will typically add its own jitter and wander as well.

Errors result from

- Timing and sampling glitches
- Impulse noise (transients)
  - Single bit (1 turning into 0 or vice versa)
- Error bursts (clusters of bit errors)
- Excessive attenuation due to lack of repeaters on long distances
  - Loss-of-signal or distorted waveforms
- Equipment failures
  - Interruptions, garbled data, distorted waveforms
- Temperature variations (wander)
- GEO satellite clock (wander)

Communications Error Definitions
The telecommunication industry generally measures performance to comply with a performance standard. Two such standards are detailed below. As no network would expect to operate with a constant bit error rate, nor would this be desirable, network performance is measured in Errored Seconds, Severly Errored Seconds and, for G.821, Degraded Minutes. These are defined as follows:

Errored Second – A second where any bit error is detected. This could be a single bit error giving a BER as low as 10^{-6}.

Severely Errored Second – Any second where more than 320 blocks out of 333 have detected errors. This can be a BER at low as 2×10^{-6} (320 bits out of 15444000) or as high as 9.6×10^{-1} (320x193x24 bits out of 15444000). Note that a BER worse than 5×10^{-1} are of no consequence as random data is 5×10^{-1}.

Degraded Minutes – The network should operate at 98.5% of 1 minute periods with BER > 10^{-6}.

Performance Standard – Accunet T1.5
Accunet T1.5 was the first industry standard commercially available T1 service. It was provided by AT&T starting in the 1980's. It guaranteed the following level of service:

- Fewer than 45 errored seconds in any 24 hour period
- Fewer than 6 severely errored seconds in any 24 hour period
- Availability of 99.9% per year

Performance Standard – G.821

The ITU-T standard G.821 "Error performance of an international digital connection operating at a bit rate below the primary rate and forming part of an integrated services network" specifies:

- Errored seconds: <1.2% of 1 second intervals
- Severely errored seconds: >99.935% of 1 second periods with BER 10^{-3}
- Degraded minutes: 98.5% of 1 minute periods with BER > 10^{-6}

While these standards apply to commercially available telecommunications networks, it may provide a guideline to utility networks as well when it comes to performance.

Errored seconds (with the subset severely errored seconds) is typically measured to determine system health and availability. A 64 kbps relay interface may be able to override errored seconds without the need for resynchronization (re-establishing communication with the remote relay) if it is just a few bit errors. It is however more than likely that events with severely errored seconds will require a re-synchronization of the relay, adding to the unavailability of the protection scheme.

Jitter and Wander

Clock errors cause jitter, wander and frame slips:

- Jitter is a short term variation in timing (> 10 Hz).
- Wander is a medium term variation in timing. It is technically the same as jitter but defined as < 10 Hz.
- A slip is when complete TDM frames are dropped or repeated in order to compensate for long term differences in timing between the node and the rest of the network. The frequency of a slip depends on network design and could be once per day or once per week. Stratum 1 (the most accurate) clock gives a slip every 72 days.

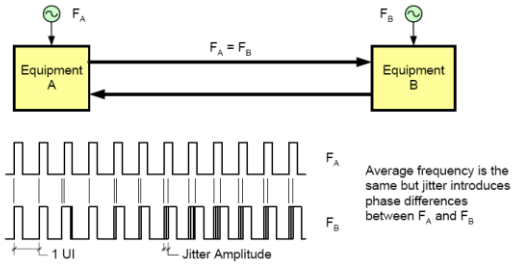

Figure 6. Jitter

Jitter is due to multiplexing processes, clock circuit design and the effects of amplitude and phase noise. Wander is due to pointer justifications (movement) in synchronous optical networks (SONET) due to the inherent architecture and made worse by temperature variations. In the case of geostationary communications satellite circuits, wander is due to the motion of the satellite in its orbit. Very low rate wander in terrestrial communications facilities is caused by the expansion and contraction of cables from diurnal and seasonal heating and cooling effects and by the varying propagation conditions along microwave radio paths. Wander, particularly very low rate wander, is impossible to eliminate.

T1 networks are designed to be primarily synchronous networks. That is, data clocked in at one point in the network has a fixed timing relationship to the point in the network at which the data is clocked out. Technically, this means that the speeds at both points are the same, and there is a fixed frequency relationship between the clocks which strobe the data in and out. This condition is usually referred to as "frequency locked".

Synchronization of the network at the DS1 level is achieved by framing the data streams and frequency locking the node and network clocks. Loss of synchronization or unlocked clocks results in frame slips. A "frame slip" is a condition in which framing is momentarily lost, as well as network timing information, typically resulting in data loss. Frame slips are typically more troublesome to protective relays than bit errors as the resultant data is now misaligned and the relay

must re-establish frame synchronization before being able to receive any data.

Effect of Frame Slips on Current Differential Relays

Referring back to the table with comparison of voice, data and relaying data characteristics, the effect of a frame slip on a relay is most likely exhibiting the same effects as for "voiceband data" in the table below.
The result is loss of the pilot channel for a period of time, and a subsequent re-synchronization. Many relays infer that a frame slip may have resulted from a communications re-routing and reconfirm the latency before re-enabling tripping. This can result in a very long outage from an instantaneous frame slip.

Table II. Effect of frame slips

Service	Effects
Voice	Audible click
Encrypted data	Loss of transmission, retransmission of encryption key
Facsimile	Corruption of 8-10 scan lines or dropped call
Video	Picture outage or freeze-frame for several seconds
Digital Data	Block retransmission; deletion or repetition of data
Voiceband Data	Carrier drop-out or error burst lasting 10 ms to 1.5 seconds

Current Differential Relaying Basics

Current differential relaying is a method of extending the benefits of differential protection as applied to transformers, buses or generators to the protection of transmission lines. Comparing current flowing into a line to the current flowing out of the same line allows for a simple protection scheme with high sensitivity and high speed simultaneous tripping of both line terminals.

The differential current can be measured with different methods:

- Magnitude comparison
- Phase comparison
- Phasor comparison (magnitude and angle)
- Charge comparison
- Combinations of the above

Regardless of the method used, all line differential relays operate on a difference in current into the line compared to the current out of the line.

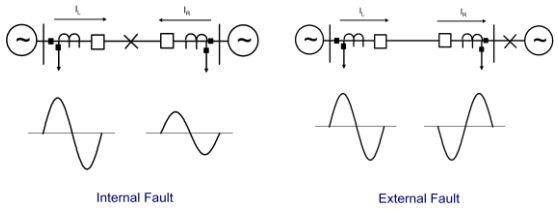

Figure 7. Current differential relaying

For an internal fault, the current will flow into the line from both line terminals, with the polarity of the current transformers. The local current I_L will be practically in phase with the remote current I_R. A small phase difference between the two currents is caused by different source angles at the local and remote end. For an external fault, the current will flow into the line in one terminal and out of the other. The local current I_L will be 180 degrees out of phase with the remote current I_R and they will be of equal magnitude.

To determine if a fault is on the line or outside the protected line section, it would be possible to examine just the differential current. For an internal fault:

$$\left|\bar{I}_L + \bar{I}_R\right| \neq 0 \qquad (1)$$

For an external fault:

$$\left|\bar{I}_L + \bar{I}_R\right| = 0 \qquad (2)$$

This simple comparison makes a differential relay very attractive for line protection as it provides a high degree of sensitivity for internal faults combined with high security for external faults.

Common to all line differential relays is the need of a reliable communications channel. A remote quantity containing the current information needs to be transferred to the local end for comparison to the local current. The quantities to be compared have to be time-coincident and the magnitude and angle information of the remote current must be preserved.

Current Differential Relay Data Frames

Different designs use different data communication frames to exchange current information between the relays at the two (or three) line ends. Not only does the way the current data is represented differ, but also the way this data is packaged into data frames for transport over the communication link.

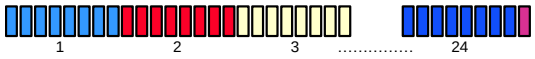

Figure 8. T1 data frame

Most current differential relays use a 64 kbps communications interface even though designs with higher bandwidths (n x 64 kbps) have recently emerged on the market. The 64 kbps data is either delivered intact over a dedicated fiber or multiplexed and ending up in a TDM channel slot, 8 bits at a time. The length of the data frame varies greatly between different relay designs. It can be as low as 15 bits or as high as 200 or 400 bits. The frequency of current data transmission also varies; from once per cycle to four times per cycle.

A 64 kbps channel provides 64,000 bits per second or 1066 bits per 60 Hz cycle. A rule of thumb is that not more than 80% of the available bandwidth is usable for payload as the rest is occupied with framing bits and error detection bits. This means that the approximate upper limit is 850 bits per cycle. A relay that sends data four times per cycle is thus limited to a data frame length of around 200 bits. Two times per cycle would allow 400 bits.

Bit errors are often detected by invalid CRC (Cyclic Redundancy Check) and cause loss of data frames. A current differential relay cannot execute its algorithm without data from the remote line end and has to wait until a valid frame is received. Until that happens, the relay will be unavailable for its primary protection function. The unavailability is longer, the longer the data frame is. For example, a 200 bit frame takes 3.125 ms which means that the relay has to wait at least this long until a new valid frame is received. A 400 bit frame takes 6.250 ms, and consequently, the unavailability due to bit errors is twice as long.

There are also relays that use a different error detection method, for example two-out-of-three voting. In this scheme, each data frame is sent three times and the receiving end accepts the message if two of the three frames are identical. This method presents an advantage in that it can override a single bit error without losing a data frame.

Another disadvantage with long frames is that there is a higher likelihood for a single bit error to disqualify a frame due to higher occupation of the available bandwidth.
A relay with a short data frame has a better chance of not having relaying algorithms degraded in noisy conditions. A bit error rate of 1/1000 is considered very high. If the frame is significantly short, <100 bits, most frames will be error free. If longer frames are used the probability is higher that the majority of frames will contain errors and be discarded. Also, shorter frames contain less information and therefore are usually of less consequence when lost. Losing a couple of bytes of data can be ridden through much easier than losing a frame with a hundred bytes.

The relay scheme unavailability also has to take into account the time it takes for the current differential relay to restore communications (re-synchronize) following a channel interruption. This time could be 30 ms or 20 seconds depending on relay design.

Typically, the substation devices are connected to a T1 multiplexer which in its turn connect to a SONET multiplexer.

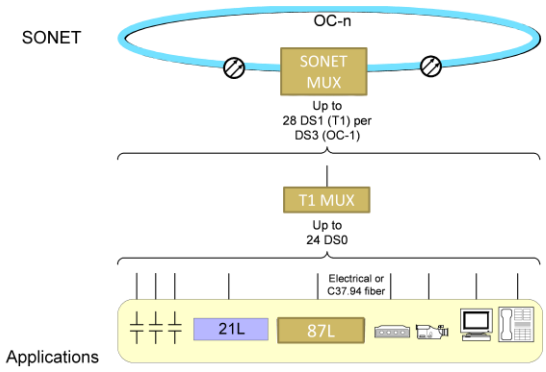

Figure 9. Protective relaying communications path

Following SONET standards, a link loss resulting in a switch-over to a healthy path takes place in less than 60 ms (10 ms detection time plus 50 ms switching time). A substation T1 multiplexer optimized for protective relaying can use "fast reframe" techniques that switches to a healthy path within a few milliseconds. Telecom T1 channel banks typically required much longer time.

Most current differential relays would react to a switching operation by having to re-synchronize the communication between the two relays. This is to ensure that the data is time-aligned properly and includes address checking to verify that the right pair of relays is communicating with each other. Depending on relay design, this re-synchronization can take 40 ms or 5 to 10 seconds and the differential protection scheme is out of service during this time period. Clearly, for relays having a long re-synchronization time, interruptions and switching operations on the communications network are going to have a considerable impact on the availability of the protection scheme.

Bit Error Testing of Current Differential Relays

The previous sections have dealt with theory. It would be of interest to test current differential relays that are subjected to bit errors typically present over a network, and record the resulting availability.

The make and model of these relays are of no importance; they are all commercial products that are in wide use. The sole purpose of the test was not to compare brands but to answer the questions:

- What requirements with regards to BER does current differential relaying put on telecommunications networks used for the pilot channel?
- Should relay bit error rate testing be part of relay evaluation testing?

Three pairs of different current differential relays were set up as follows:

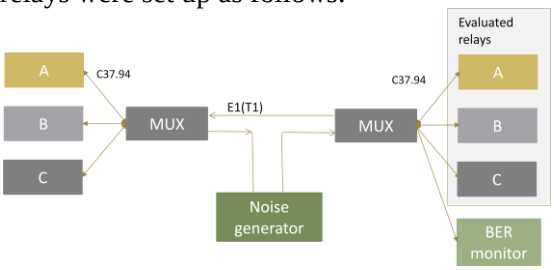

Figure 10. Bit Error Rate test

The three relays were equipped with C.37.94 interfaces, each of them using one 64 kbps channel over the E1/T1 substation multiplexer. The T1 communications link was routed via a noise generator, injecting an adjustable level of noise onto the T1 circuit. Fast reframe technology was used in the multiplexer which means that re-frame of the E1/T1 link was achieved in less than one ms for loss of channel conditions caused by excessive bit errors.

The average BER was measured with a BER tester, connected to a channel card on the multiplexer.

Testing was made with different average Bit Error Rates and the resulting availability recorded in the table below.

Each test was run for a period of 15 minutes but as the results were consistent over this time, only the last 2 minutes were used to calculate availability. The identical time period was used for recording the results for all three sets of relays.

Availability was determined by examining the relay SOE record on the RX (receiving) end of the link subjected to noise. SOE was triggered by programmable signals in the relays that reflect whether the current differential protection is active or not.

Table III. Relay availability during Bit Errors

BER	*Availability (%)*		
	Relay A	Relay B	Relay C
10^{-6}	100	100	100
10^{-5}	72	96	100
10^{-4}	46	0	100
10^{-3}	50	0	51

The tests show that all relay performed well for BER 10^{-6}.

Typical error rates for copper and optical transmissions are in the range 10^{-10} to 10^{-14} therefore a requirement of $<10^{-6}$ should be easily met during quiescent conditions.
(As a comparison, cellular wireless networks can have BER as low as 10^{-3} to 10^{-6}.)

Differences between the relay designs started to show up for BER 10^{-5} and worse. The differences are likely due to differences in the design when it comes to the time required for recovery following a channel interruption. This difference may have been less obvious if the test had simulated the errored seconds or severly errored seconds used in the performance standard rather than providing a constant BER. A short burst of noise followed by a relatively longer period of good channel would have reduced the impact of re-synchronization time on availability.

However, the results show that there is not much difference in the availability for BER that could be expected on a network which leads to the conclusion that network design is the dominating factor when it comes to protective relay scheme availability. Low availability over an extensive network is a factor of network design, equipment chosen, and maintenance practices.
In regards to the second question; whether BER testing should be part of relay evaluation, there is more than one answer. If the application will be over dedicated fiber, a back-to-back test without noise injection is perfectly sufficient. If the application will be over a network, perhaps it is a good idea to verify that the relay will handle BER expected over that network, or get assurance from the manufacturer that this is the case.

Conclusions

While a dedicated fiber pair is the preferred communications link for a relay engineer, this alternative is not always an option. There may be no direct point-to-point fiber available and it is also an inefficient use of capital investment as the fiber can carry much more data than just relaying, if multiplexed.

However, to ensure a reliable relay scheme with high availability, the network design needs to take into account and prioritize meeting the requirements by the relay scheme. This includes avoiding unnecessary intermediate devices, minimizing buffering, and minimizing switching operations on the network.

The tests performed verified that the relay interface is not the critical factor for availability. While the tested relays were all of different design, the performance was comparable.
The tested relays all performed well with a BER exceeding what can be expected on a digital communications network. However, this does not guarantee that all relays on the market will do the same.

Network errors cannot be eliminated but with a proper network design, the effect on relays can be minimized.

Biographies

Solveig M. Ward

Solveig received her M.S.E.E. from the Royal Institute of Technology, Sweden in 1977. The same year she joined ABB Relays. She has held many positions in Marketing, Application, and Product Management. Assignments include a six-month period in Montreal, Canada and two years in Mexico. When Ms. Ward returned to Sweden, she was responsible for the application aspects in the development of a numerical distance protection relay and in charge of marketing the product. After transferring to ABB in the US 1992, she was involved in numerical distance protection application design, and was Product Manager for ABB's line of current differential and phase comparison relays. In 2002, Solveig joined RFL Electronics Inc. as Director of Product Marketing where she was involved in the development of new relay and communication products. Solveig joined Quanta Technology in 2010 as Principal Advisor in the Automation and Protection group where she is a team leader for Automation. She is presently working on projects involving IEC 61850, communications and Cyber Security.

Solveig has written, co-authored and presented several technical papers at Protective Relaying Conferences. She is a senior member of IEEE and holds one patent, "High Speed Single Pole Trip Logic". Solveig serves as the Chair of the PSRC System Protection Subcommittee.

William G. Higinbotham

Bill has been Vice President of RFL Electronics Inc. Research and Development Engineering group since 1994. He joined RFL 21 years ago as a senior design engineer. In this period he has been involved in the development of numerous RFL products in the areas of communications and protection. Bill is active in the IEEE Power Systems Relaying Committee and has written and coauthored a number of technical papers. He holds one patent related to multiplexer resynchronization times. Bill received his BSEE from Rutgers University in 1984 and worked in the biomedical engineering field for 5 years.