Back to Table of Contents

DIGITAL COMMUNICATIONS FOR POWER SYSTEM PROTECTION: SECURITY, AVAILABILITY, AND SPEED

EDMUND SCHWEITZER III
KEN BEHRENDT
TONY LEE

SCHWEITZER ENGINEERING LABORATORIES, INC.
PULLMAN, WASHINGTON USA

Presented before the

25th ANNUAL
WESTERN PROTECTIVE RELAY CONFERENCE
SPOKANE, WASHINGTON
OCTOBER 13-15, 1998

DIGITAL COMMUNICATIONS FOR POWER SYSTEM PROTECTION: SECURITY, AVAILABILITY, AND SPEED

Edmund Schweitzer III, Ken Behrendt, and Tony Lee
Schweitzer Engineering Laboratories, Inc.
Pullman, WA USA

INTRODUCTION

New channels and digital techniques in communications provide opportunities to advance the speed, security, dependability, and sensitivity of protection – while simultaneously reducing the costs associated with using communications. Lower communication costs mean more opportunities to benefit from pilot protection. The net result is a higher quality of electric power delivered for each dollar invested.

A classical pilot communication scheme is shown in Figure 1, and a direct digital-to-digital scheme is shown in Figure 2. Clearly, the direct digital system is simpler. In this paper, we will show that direct digital communications economically provide several bits in each direction – and these extra bits lead to simpler, more sensitive, and more flexible schemes.

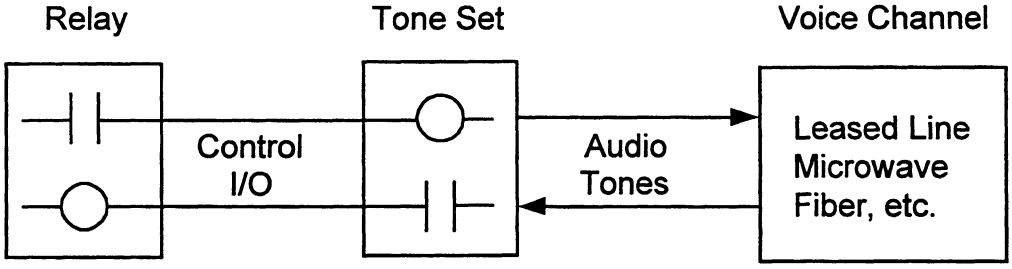

Figure 1: Frequency-Shift Audio Keying Over Voice Channel

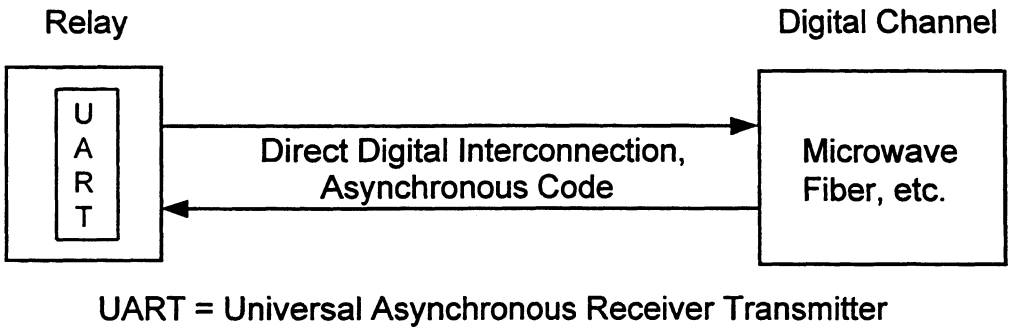

Figure 2: Direct Digital Signaling Over Asynchronous Channel

WHAT IS THE CAPACITY OF A CHANNEL TO COMMUNICATE?

How much information can be sent, theoretically and practically, through a given channel while still maintaining acceptable reliability?

In 1948, Claude E. Shannon [1] mathematically formalized a theoretical limit. His theory was that information can be reliably transmitted over a noisy channel if the data transmission rate is sufficiently low. If the relative noise increases, the maximum reliable transmission rate decreases. For instance, a channel with bandwidth W and received noise power N, can transmit information at rate C with arbitrarily high dependability, as long as the average signal power P satisfies:

$$C = W \log_2\left(\frac{P}{N} + 1\right).$$

The ratio P/N is the signal to noise ratio (SNR). The base of the logarithm depends on how we measure C. When C is measured in bits per second, the logarithm is base two. In general, the logarithm base is the number of symbols in the alphabet to be transmitted.

Quieter channels lend themselves to faster data transmission. Conversely, faster data transmission requires a quieter channel.

If SNR = P/N = 1, then the Shannon limit is C = W (bits/second), i.e. the channel capacity for reliable transmission is the channel bandwidth.

If the SNR = 20dB = 100, then C = 6.7 W (bits/second). We can see that increasing the SNR gives us an opportunity to reliably transmit more data, faster.

For example, contrast the channel requirements for two different transmission rates. A frequency shift keyed (FSK) audio tone transmission over a voice channel might carry a single bit of information, such as a permissive trip signal. Suppose we require transmission to occur reliably in 20 ms, so the required data transmission rate is 1 bit / 20 ms = 50 bits/second. Also assume the receiver filter has a bandwidth of 300 Hz. According to Shannon, the received signal, after filtering, must have an SNR of greater than about 0.1. Before filtering, the SNR on a 3 kHz channel must be greater than about 0.01, assuming white noise.

A 9600 bits/second data stream transmitted over a 3 kHz channel requires an SNR of at least about 10, according to Shannon's work. Therefore, the voice grade channel must be about 10/0.01 or 1000 times quieter (assuming the same transmit power and modulation techniques) to reliably transfer data at 9600 baud.

In practice, it is difficult to approach Shannon's limits. The two examples cited above, when implemented using present technology, both need about ten times better SNR than Shannon's limit.

Both Shannon's prediction and practical experience show that when we have a better channel, we can send more information per unit bandwidth. The quality of many digital channels is excellent, and opens the door to new digital techniques in protection.

The communications engineer uses encoding and modulation techniques to approach Shannon's limit, and to balance data rate with reliability within the context of a given application.

ENCODING EIGHT BITS FOR DIGITAL TRANSMISSION

We foresee a great future for sharing a handful of bits directly from one relay to another, over an array of digital channels of moderate capacity. Pilot protection, control, adaptive relaying, monitoring, and breaker-failure are some examples.

Our starting point was eight bits of information in a message with enough redundancy to meet protection-security requirements, yet efficient enough to be useful at data rates from several kilobaud and up.

Security

Communications are secure when the receiving end reliably detects whether the received information differs from the transmitted information.

The standard IEC 834-1 [2] contains recommendations for blocking, permissive, and direct tripping pilot schemes, in terms of their susceptibility to noise bursts. The short table below gives the expected minimum number of noise bursts required to produce an undesirable output.

Scheme Type	Security (bursts/undetected error)
Blocking	10^4
Permissive Tripping	10^7
Direct Tripping	10^8

To help detect noise bursts, we can add some redundant information to the transmitted message. Shannon gives a formal definition of redundancy:

Redundancy = Total Bits Transmitted - Information

For example, if we transmit a total of 10 bits, and there are eight bits of information, then the redundancy is 10 - 8 = 2 bits.

Redundancy is necessary, but not sufficient for security. One of the objectives of encoding is to make each of the distinct messages (e.g., for eight bits there are 256 different messages) as different as possible from the rest. The quantitative measure of this difference is the Hamming distance. It is defined as the minimum number of bits that could be corrupted in one distinct message, which would result in a different distinct valid message.

The simplest form of redundancy which increases Hamming distance is repetition. Consider a single bit of information (permissive trip for instance) which must be received by the remote relay from the local relay. If the local relay transmits only that bit, the remote relay cannot detect whether the bit has been corrupted. The remote relay receives a one or a zero, and has no indication if the received value is the same as the transmitted value.

Now assume that the local relay transmits the bit of information twice. The receiving relay compares the two bits. If they are the same, the receiving relay assumes they are correct and accepts the bit of information. But, if they differ, the receiving relay discards the information.

When we inject noise bursts onto the channel, we generate messages almost at random, because the bit error rate is so high. There are two valid messages (00 and 11) and two invalid messages (10 and 01). Therefore, we expect that a randomly generated message will be accepted by the receiving relay half the time, or once per two noise bursts. Since that is a long way from 10^7, add more redundancy.

If the transmitting relay adds another bit of redundancy, then there are still only two valid messages (000 and 111), but there are now six invalid messages (001, 010, 011, 100 101, and 110). The receiving relay will accept a randomly generated message two out of eight times, or one out of four, on average.

Every time we add a bit of redundancy, we cut the probability of accepting a randomly generated message by half. Thus, the expected number of randomly generated messages is 2^n per undetected error, where n is the number of redundant bits.

We need $\log_2 (10^7)$, or about 23.3, bits of redundancy to get 10^7 security. Our 1 bit of information, plus 24 redundant bits, yields a 25-bit message.

This method does not make very good use of our channel however, because we must transmit 25 bits to securely communicate a single bit of information. Now that we have the required security, we can add more information bits with no loss of redundancy.

Suppose we decide to repeat eight bits of information three times, and add some channel framing bits for 36 bits total. Four of these channel framing bits do not count as redundant bits, so there are 36 - 4 - 8 = 24 redundant bits. We have increased our information transmission capability by a factor of eight over the original single bit, decreased the rate of transmission by 1 - 36/25 = 44%, and maintained $2^{24} = 1.7 \times 10^7$ security to randomly generated messages.

Dependability

Suppose we want to transmit eight data bits with 10^7 security as described above. We have shown that the message must consist of at least the eight data bits plus 24 redundant bits, plus some non-redundant channel framing bits (36 bits total).

Assume that channel errors occur at some average random bit error rate (BER). The receiver rejects an entire message of 36 bits even if just one bit is in error. In addition, after the receiver detects a bit error, it will probably require that two or more consecutive messages are received without error before using the received information. The probability that a message will not be accepted is about $2 \bullet k \bullet BER$, where k is the number of bits in the message. This approximation holds for $2 \bullet k \bullet BER < 0.1$. Therefore, we expect the average unavailability to be about 72 times the channel bit error rate. By increasing the redundancy from one to 24 bits we have increased security from 2 to 10^7. Simultaneously, we increased unavailability by a factor of 72. If we started with a channel with BER of 10^{-6} the unavailability is now 72×10^{-6}. This demonstrates a trade-off between security and dependability: increasing security by a factor of 10^6 decreased dependability by a factor of 72.

Speed

Again consider the 36-bit message developed earlier. Compared to the two-bit message with a single redundant bit, it takes 18 times as long to transmit/receive. At 9600 bits/second, it takes

about 3.6 ms to transmit/receive the 36-bit message, compared to 0.2 ms for the 2 bit message. Therefore, increasing security decreases speed somewhat.

Adaptability

The 36-bit message developed above gives 10^7 security, and, when coupled with the proper digital channel, still yields high speed and excellent dependability. Suppose we use one of the eight data bits for a block-trip signal, and another of the bits for a remote-control direct-trip signal. The security afforded by the 36-bit protocol is sufficient for blocking schemes. However, IEC 834-1 recommends ten times better security for direct tripping. We want to increase the security of the direct trip signal without affecting the speed or availability of the blocking signal.

This is easily accomplished with a pickup security counter on the direct trip bit. For example, a count of two requires reception of two successive 36-bit messages with the direct trip bit set before updating the direct trip bit in the receiving relay. If we return to the test prescribed by IEC 834-1, we would still expect to inject 10,000,000 bursts of noise, on average, to get one corrupted message which is incorrectly accepted by the receiving relay. However, to perturb the direct-trip bit qualified by a two-count security counter, the very next message must also be acceptable, and must have the same direct trip value. This also happens about one in 10,000,000 times. So the probability of a false trip in response to noise bursts is about $(10^7)^2$ or 10^{14}. This is six orders of magnitude more secure than the IEC 834-1 recommendation for direct tripping.

Remember we must trade off speed and/or availability to gain security. Here, the direct trip signal is delayed by one additional message-time, and the unavailability is roughly doubled.

Practical Implementation

We implemented the 36-bit code described above. The 36-bit message is transmitted at 19,200 kbits/s in 36/19,200 = 1.875 ms. Allowing for 2 ms of latency, plus 2 ms for processing time in the receiving relay, gives a total of about 6 ms from the time the transmitting relay makes a decision to when the receiving relay has makes a decision influenced by the transmitting relay.

The delays for a tone set between two relays are much longer:

2 ms output contact + 12 ms tone set + 2 ms latency + 2 ms processing = 18 ms.

Thus, the direct digital communications gives us eight times the data with one-third the delay, at far less cost and complexity.

To test the protocol security, we injected 200 ms long white noise bursts onto a direct copper connection between relays. We set the transmitting relay to transmit a known set of eight bits, and we set the receiving relay to trigger an event report upon reception of anything but that known pattern.

The receiving relay triggered the first event report after 7 million noise bursts. We terminated the test after 20 million noise bursts (and nearly 50 days) with still only one undetected error.

Applicability

Since the protocol described above is a simple serial bit stream, it is compatible with many channels and many types of data communications equipment.

Channel Performance Monitoring

Digital communications provide opportunities for performance monitoring, so the quality can be assessed, and problems can be quickly detected and remedied.

One channel monitor tallies the time the received data are corrupted or absent, and normalizes that time to the total elapsed time. This directly measures the unavailability of the communications. Although unavailability is a useful long-term measurement, it hides long but infrequent channel disturbances. For example, suppose a channel monitor is set to alarm when the unavailability exceeds 500×10^{-6}. If that channel is error-free for one year and then the channel is completely lost, the unavailability monitor will not alarm until four hours later.

A second monitor can be used to alarm when the channel is not available for a certain continuous time, say one second.

The unavailability alarm responds to gradual degradations in bit-error rate. The duration alarm responds more quickly to outright channel failures.

A sample report from such a monitor follows. It reports the 256 most-recent errors, the average unavailability for the time of the report, and the longest-duration channel outage.

```
Summary for Channel A

For 06/19/98 15:43:48.887 to 07/30/98 10:13:11.925

    Total failures     14              Last error  Re-sync
    Relay disabled      1
    Data error          4              Longest failure  0 00:00:41.352
    Re-sync             4
    Underrun            1              Unavailability   0.000015
    Overrun             0
    Parity error        3
    Framing error       1

    START    START         END      END
#   DATE     TIME          DATE     TIME          DURATION      EVENT
1   07/10/98 11:19:14.769  07/10/98 11:19:24.419  00:00:09.650  Re-sync
2   07/09/98 11:48:13.572  07/09/98 11:48:14.126  00:00:00.554  Underrun
3   07/09/98 11:48:12.710  07/09/98 11:48:13.481  00:00:00.770  Re-sync
4   07/09/98 10:38:32.062  07/09/98 10:39:13.414  00:00:41.352  Parity error
5   07/07/98 09:33:35.389  07/07/98 09:33:35.419  00:00:00.029  Re-sync
6   07/07/98 09:21:44.183  07/07/98 09:21:44.229  00:00:00.045  Parity error
7   07/07/98 09:21:44.087  07/07/98 09:21:44.154  00:00:00.066  Data error
8   07/07/98 09:21:36.077  07/07/98 09:21:36.127  00:00:00.049  Data error
9   07/07/98 09:21:33.727  07/07/98 09:21:33.777  00:00:00.050  Data error
10  06/29/98 09:19:12.075  06/29/98 09:19:12.120  00:00:00.044  Framing error
11  06/26/98 15:04:28.653  06/26/98 15:04:28.701  00:00:00.047  Data error
12  06/26/98 15:01:40.209  06/26/98 15:01:40.243  00:00:00.033  Re-sync
13  06/26/98 15:00:27.803  06/26/98 15:00:27.845  00:00:00.041  Parity error
14  06/19/98 15:43:48.887  06/19/98 15:43:48.887  00:00:00.000  Relay disabled

=>
```

This specific channel is a digital leased line running at 56 kbaud. The relay detected 14 total errors that resulted in an average unavailability of 15×10^{-6}. The longest channel outage was 41.352 seconds.

Notice that the errors are grouped in clumps. The report was cleared on 6/19, and the circuit experienced no errors until 6/26. On 6/26 there were three errors in four minutes. The circuit was then perfect for three days, until a single error occurred on 6/29. There were no more errors for over a week, then there were four errors in eleven seconds, followed by another error twelve minutes later. After two days without error, there was a span of 41 seconds where the relay did not receive an acceptable message. This underscores the additional value of a continuous outage monitor, because the unavailability monitor would not have alarmed for this outage unless it was set as low as 30×10^{-6}. On the next day another extended outage of over nine seconds occurred. The channel was error-free for the next 20 days, from 7/10 to 7/30, when the report was downloaded from the relay.

Later, we will discuss error-seconds per day as a measure of quality of the leased lines. Performance monitoring provides the quantitative feedback needed to maintain and improve the quality of communications. Relay event reporting provides an additional perspective on the performance of communications for every event, because the communicated bits are reported as additional I/O points. For example, communications disruptions during faults are easily observed, should they occur.

CHANNELS

This section compares some communications channels that might be used in pilot and control schemes.

Dedicated Fiber

Perhaps the ultimate digital channel in terms of dependability, security, speed, and simplicity is dedicated fiber optics. Low-cost fiber-optic modems make dedicated fiber channels even more attractive. Often, the modems can be powered by the relay, eliminating the cost and loss of availability involved in separate power sources. Some modems also plug directly onto the digital relay [3], which eliminates a metallic cable. Eliminating the cable and the external power source removes "antennas" for possible EMI susceptibility.

When the communications path is short, the cost of the fiber is not very significant. On longer paths, multiplexers may be considered, to increase the amount of data communicated over a fiber-pair. However, the relatively small incremental cost of adding and using one fiber-pair for protection alone is probably justified by the increase in simplicity and availability that a dedicated fiber scheme offers. Furthermore, the small incremental cost is partially offset by the very low cost of the simple direct-connected fiber-optic modems.

Bit errors are extremely rare on most fiber-optic links. If the link is available, then it is near-perfect, because the fiber medium is unaffected by the RFI, EMI, ground-potential rise, weather, and so on.

The receiver amplifier is the major source of noise in the system – and that noise source is very small compared to the large signals used in simple, practical systems. The received signal

strength is the transmit power minus attenuation. Attenuation in decibels is proportional to fiber length plus some loss for each connection or splice. A system designer usually fixes the transmit power at some level, specifies a power margin, and then guarantees some allowable fiber length at some maximum bit-error-rate or BER. If we use less fiber or fewer connections, there is more signal power at the receiver, and the BER decreases.

Suppose the system designer chose a maximum allowable BER of 10^{-9} at some maximum allowable fiber length. Figure 3 was adapted from [4]. It shows that decreasing the fiber length by 40% in such a system (to 0.6 per unit), decreased the BER from 10^{-9} to about 10^{-23}, a decrease of 14 orders of magnitude! Random bit errors cease to be an important source of unavailability.

Unavailability then becomes dominated by other factors such as fiber breaks, misapplications, etc. Therefore, well-designed fiber-optic communications systems will result in long periods of error-free performance separated by complete outages caused by human factors or equipment failure.

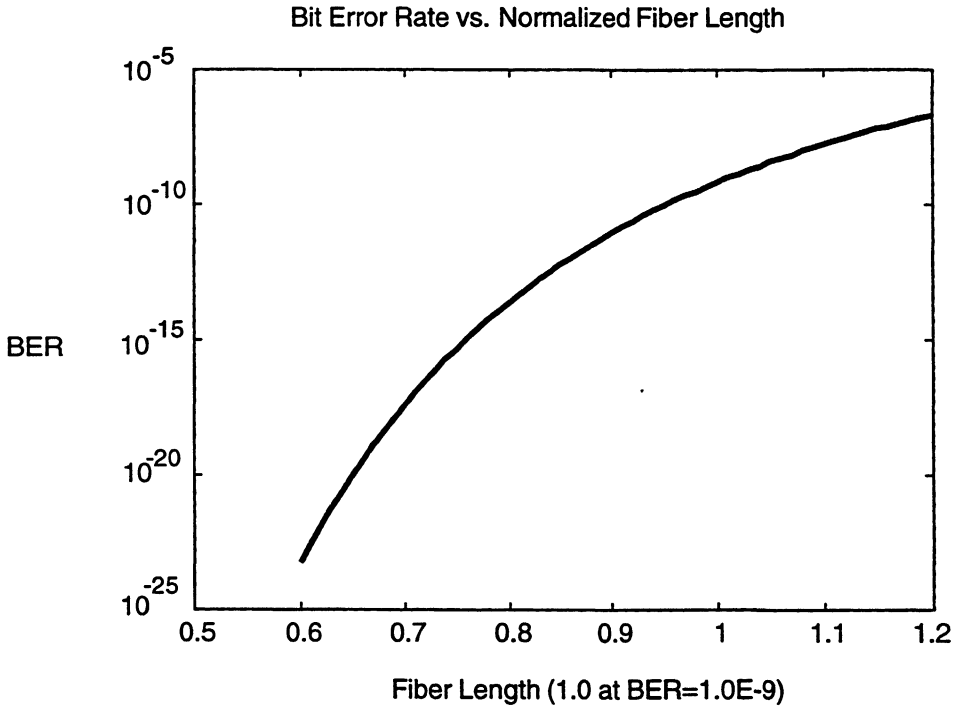

Figure 3: Conservative Designs Yield Near-Zero BER on Fiber Links

Does this mean we can ignore communications security? No. Consider what happens to a fiber-optic receiver when a user disconnects one end of the fiber while the link is in service. The disconnection process is slow compared to most fiber-optic data transfer rates. As the user draws the fiber away from the receiver, attenuation increases and ambient light begins to flood the receiver. This causes the bit error rate to increase until the received bit stream is essentially all noise. The receiving device must recognize this noise and reject the corrupted data, or a misoperation may result. The security built into the 36-bit message described earlier is sufficient to ensure less than a 10^{-7} chance that disconnecting the fiber could cause an undetected error. A security counter of just two virtually eliminates the risk, even for direct tripping.

Can direct fiber channels be affected by faults? There is some risk that the physical event which breaks the fiber could cause the fault, such as a tower collapse or static wire failure where the fibers are in the static wire. Tornadoes, ice loading, or an airplane collision are also possibilities. Even then, there is some chance the message will get through to permit the scheme to work, after the fault occurs and before the channel is destroyed.

When a dedicated fiber is closely associated with the power line right-of-way, the probability that an external fault will cause a communications disturbance is negligible.

Multiplexed Fiber

Fiber-optic multiplexers combine many relatively slow digital and analog channels into one wideband light signal. The multiplexer, therefore, makes efficient use of bandwidth in the fiber. A direct digital connection between the relay and the multiplexer is more reliable and economical than interfacing through conventional relay contacts, then a tone set, and into an analog channel on the multiplexer. The multiplexer adds a level of complexity, which can be avoided by the simple dedicated fiber approach discussed earlier.

Fiber-Optic Networks

Wide-area networks, such as SONET, move large quantities of data at high speed. Many such networks consist of self-healing rings.

Since the self-heal time is long compared to expected protective relay tripping times, we must still be concerned with correlation between faults and communications problems.

There is a tradeoff between long-term availability and short-term dependability. The ring self-heals so that communications are rarely totally lost. However, a failure anywhere in the network results in a short communications loss. Power Networking [5] describes a cascaded ring topology which reduces the exposure to these short interruptions.

While the ring is self-healing, the terminal equipment is generally not. Thus the terminal equipment, and possibly other points, must be considered as possible single points of failure – even though we have a self-healing ring.

Multiplexed Microwave

Microwave systems have gone digital, too – opening new opportunities for direct relay-to-relay communications. (Later we describe a low-delay modem, which can be used to transfer digital information through analog microwave channels, with the quality required for pilot protection.)

Microwave equipment failures include multiplexers, radio gear, antenna pointing errors, cabling, etc. Microwave communications are fairly immune to power system faults. In general, the likelihood of a communication failure for an internal fault is not much different from the likelihood of a failure for an external fault.

Narrow-Band UHF Radio

Dedicated radios have been used for pilot channels. Reference 1 describes how a 960 MHz radio link was used in a POTT scheme. The radio was purchased with a single on-off-keyed tone

interface between the radio and the relay contact I/O. The security of this scheme comes from the inherent security of POTT schemes.

Narrow-band digital radios permit the use of direct digital communications. The performance would be similar to that of the microwave system given earlier, but may be more reliable than a channel in a microwave system, because of lower complexity.

Radio channels are relatively immune to interference from faults. One possible source of interference is the power wiring to the radio. Radio channels might also be disrupted by antenna-pointing errors and severe weather.

Radio and fiber channels can be highly complementary. Mechanical damage that might disrupt a fiber channel is generally unlikely to interfere with a radio channel, and vice versa. However, it is possible to conceive of events that would destroy both communications channels and cause a fault. For example, suppose a radio tower collapses during an earthquake or windstorm, and falls through the static wire with the fiber in it, and then causes a line fault!

Spread-Spectrum Radio

Spread-spectrum techniques have been broadly applied in radar systems to increase the energy in a radar pulse, while maintaining and enhancing target resolution. Spread-spectrum communications are used in military applications for the advantages of communications security, interference immunity, low probability of detection, and difficulty in jamming. Commercial uses of spread-spectrum radio have been growing, ever since the Federal Communications Commission permitted license-free operation under certain conditions. For power system protection, the advantages of spread-spectrum radio channels are immunity to interference, freedom from licensing requirements, and low cost.

Signals may be spread in the frequency domain by several methods. For example, frequency hopping, either slow or fast compared to the information rate, spreads the signal over the spectrum covered by dozens of discrete frequencies occupied sequentially in time, in a pseudo-random sequence. Direct-sequence spread spectrum systems, on the other hand, multiply the information bit stream by a much faster pseudo-random binary sequence. The bandwidth is expanded by the fast rate of the pseudo-random sequence.

The processes of spreading, despreading, synchronization, and forward error correction (FEC) take some time, and, depending on the scheme, may be too slow for teleprotection. Most presently available radios also rely on the data terminal equipment (DTE) to negotiate a half-duplex channel. However, at least one model automatically switches its half-duplex channel rapidly enough to simulate full-duplex transmission at speeds as high as 19,200 baud. This same radio performs no FEC, and so has a very reasonable round-trip delay of about 2.5 power system cycles. The cost, power requirements, and performance are promising for applications all the way down to distribution voltages [7].

Digital Telephone Circuits

Telephone companies offer leased digital lines for several hundred dollars per month, and these can be used for pilot protection schemes. A CSU/DSU interfaces the protective relay to the leased line. It receives timing information from the telephone company equipment via the leased line, and passes that timing information on to the relay (for synchronous data) or synchronizes

the asynchronous data stream from the relay (for asynchronous data). It also converts the serial data received from the relay to the proper electrical levels and format.

The digital data are not modulated on the twisted pair in the traditional sense; they remain binary (actually ternary) while on the leased line. Such communications are characterized by long periods between short bursts of errors. For example, the standard AT&T Technical Reference TR 62310 [8] defines acceptable performance of a 56kbps using the concept of an error-second (ES) and a severe-error-second (SES). At 56kbps, an ES has between one and 56 bit-errors. A SES has more than 56 bit-errors. That standard allows 20 error-seconds and six severe-error-seconds per day.

If we assume each ES results in 1 second of protection scheme unavailability, then communications-assisted protection might not be available for 26 seconds per day. This is a very extreme upper limit.

Like fiber-optic links, digital leased lines are distance-sensitive. Far from the distance limit, the actual performance is significantly better than the worst-case prediction above. Never accept anything close to the worst-case scenario depicted above. We have experience with digital leased lines that produced less than one error every three days. Earlier in this paper, we presented data from a digital line that was unavailable for approximately 1 minute during a 40-day period. This is 17 times better than the worst case scenario.

Another factor affecting error rates on digital leased lines is transmit power. In AT&T Data Communications TR 62310 transmit power is restricted at 9,600 and 12,400 bits per second to 1/4 of the power allowed at other rates. At least one CSU/DSU manufacturer takes advantage of the higher allowed transmit power at data rates other than 9,600 and 12,400 bits per second. As with fiber optics, increasing the transmit power by a factor of four on a twisted pair can have a profound impact on random bit errors. Assuming a leased line takes advantage of the higher permitted power, then rates other than 9,600 are greatly preferred.

Any leased line must be galvanically isolated between the substation and the central office. This isolation prevents damage and danger when ground faults produce high voltages between the substation ground and the telephone exchange [9]. However, isolation does not guarantee that the leased line will remain operational during the fault. Ground potential rise or noise coupled from the faulted power line to the twisted pair can produce enough noise on the circuit to cause bit errors or a complete loss of signal. The analysis required to determine if a circuit will remain operational would be difficult.

Noise from faults or other sources that might not corrupt signaling by audio tones over a given twisted pair still might interfere with fast digital signaling over that same channel. (Recall that faster signaling requires a higher SNR, given the same bandwidth.)

Analog Voice-Grade Channels

Figure 4 shows a low-delay modem interface between the direct digital data from the protective relay, and a voice channel. The voice channel may be analog microwave, a leased circuit from the telephone company, a dedicated twisted pair, or something similar. Analog channels on microwave should not suffer the same degradation during power system faults as analog channels on a twisted pair or some other conductor. The challenge becomes modulating the

9600 bit per second data stream so it will be compatible with the 300 to 3,000 Hz audio band channel.

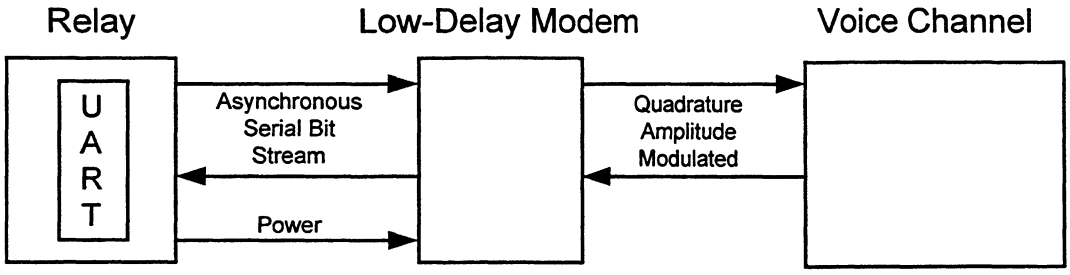

Figure 4: Direct Digital Signaling Over Voice Channel

Computer modems generally are unsuitable for protection. They are optimized for throughput, at the expense of delay. Therefore a special low-delay modem was developed. It also takes a very short time to adapt to changes in the channel (retraining), compared to computer modems [10].

PILOT PROTECTION APPLICATIONS

Direct underreaching transfer trip schemes require extremely secure communications, because there is no local confirmation that a fault exists. DUTT is very simple. The digital communications described in this paper provide much greater security than that required by IEC 834-1 for direct tripping, when the protection scheme uses a security count of two or greater. It is inherently secure from current reversals.

Permissive underreaching transfer trip does not require as secure a level of communications, because the received underreaching element is qualified by a local overreaching element. The scheme is also quite simple, and requires no current-reversal logic. Sensitivity is very similar to DUTT.

Permissive overreaching transfer trip schemes provide greater sensitivity, and have the same channel dependency as PUTT. In most cases, POTT schemes must be protected against current reversals. POTT schemes handle weak terminals, when the schemes include echo logic. If internal faults cause channel failures, then POTT schemes may not operate.

Directional comparison blocking schemes provide very similar speed and sensitivity to POTT with echo logic, yet DCB schemes do not require echo logic. DCB schemes must also be protected against current reversals. If external faults cause channel failures, then DCB schemes will overtrip. The security and practicality of DCB schemes depend on known and reasonable upper limits on element pickup times and channel delays.

Directional comparison unblocking schemes attempt to give the best of POTT and DCB. DCUB schemes only make sense when we can definitely associate a much greater likelihood of channel failures with internal faults, than with external faults. For example, DCUB might be sensible for a power-line carrier channel, or an optical-fiber in the shield wire of the protected line.

Consider some pilot-scheme possibilities, given different channels.

Fiber-Optic Ring

A good approach is POTT, with weak-infeed and open-breaker echo. DCB should not be used because: security depends directly on availability of the communications, we cannot associate channel failures with internal faults, and communication delays may depend on routing. DCUB should not be used, because we cannot associate channel failures with internal faults.

Leased Digital Line

As with the ring, it is not generally possible to associate channel failures with either internal or external faults, and delays may be variable. Therefore DCB and DCUB should be avoided.

Dedicated Fiber Optics

If the fiber and the power line share the same path, then DCUB might be used to gain some availability, with little loss in security. This is because channel failures simultaneous with faults might reasonably be associated with internal faults.

Sending multiple bits in each direction opens up some new possibilities in pilot protection.

Cross-Country Faults

Consider a double-circuit line from S to R, as shown in Figure 5. An AG fault on Line 1 simultaneously exists with a BG fault on Line 2. In a single-pole-tripping scheme, the desired action is for Line 1 to trip phase A, and Line 2 to trip phase B, so service is essentially uninterrupted between S and R. If the relays at R communicate the observed fault types, then the relays at S can trip single-pole and avoid the undesired three-pole trips for this situation. The fault type is easily communicated on three bits, or on just two, with some encoding.

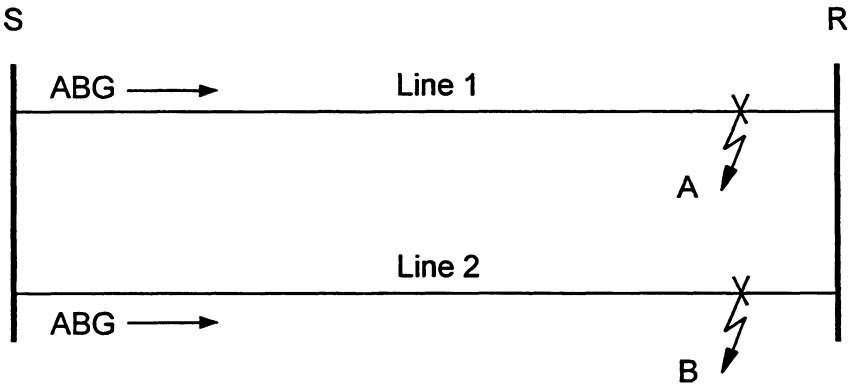

Figure 5: Communicate Fault Type for Secure Single-Pole Tripping

Remote-End Open Keying

Given a line from S to R, assume the breaker at R is open. A POTT scheme needs to inform S that R is open. Traditionally, there are two ways. One is for S to send permission to R and for R to "echo" the permission back to S based on the open-breaker status. This status is sensed by monitoring an auxiliary contact at R. Breaker S trips after two communications delays, plus

possibly an echo time-delay in terminal R. Given 16 ms for the protective relay element at S, 12 ms for each communications delay, and no echo time-delay in R, we have a total time of 40 ms from fault inception until S trips. (Another timer is used to limit the echo duration to a few cycles, so the channel does not lock up.)

A second method is for terminal R to send S a standing permission whenever the breaker at R is open. Both communications delays are eliminated, so tripping can occur in just relay-element time, e.g., 16 ms. With conventional channels, this method makes a compromise with security because guard-before-trip cannot be used. However, with direct digital communications, the security is built into the message so no loss of security occurs with hard-keying.

An alternative is shown in Figure 6, where the remote end transmits its breaker status and breaker control commands to the local end. Overreaching elements trip the local breaker, as long as the remote end is open. The philosophy here is for the remote end to send the local end the state of the breaker so that the local end can directly observe it and use it as desired, in this scheme and possibly others. If the local end receives notice of close commands from the remote end, then this scheme can be briefly delayed to avoid risk of misoperation by very sensitive elements due to pole scatter. Other advantages are that this scheme is never encumbered by current-reversal timers, and it is possible to use different overreaching elements when the breaker is open than when it is closed.

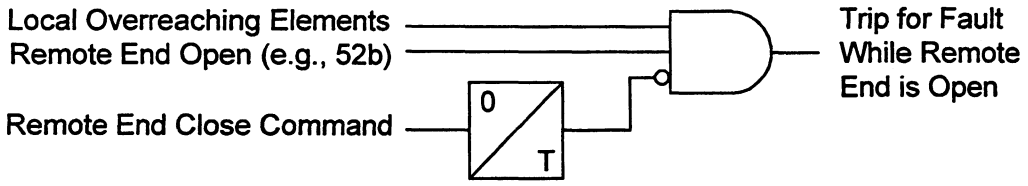

Figure 6: Line Protection for Open-Remote-End

Simplify Pilot Logic

When a single bit is available, say in a POTT scheme, a timer and additional logic are necessary to avoid latching the channels on during echo. The logic is simpler if we elect to use individual bits to transmit the status of individual relay elements, and of the breaker. Each terminal builds up its trip outputs from the locally-observed and remotely-reported relay elements without the need for feedback paths that lead to channel lock-up or other surprises. Figure 7a shows how a traditional POTT scheme avoids channel lock-up with extra logic and a timer. Figure 7b shows a simpler scheme, which does not have the risk of lock-up because there is no feedback path.

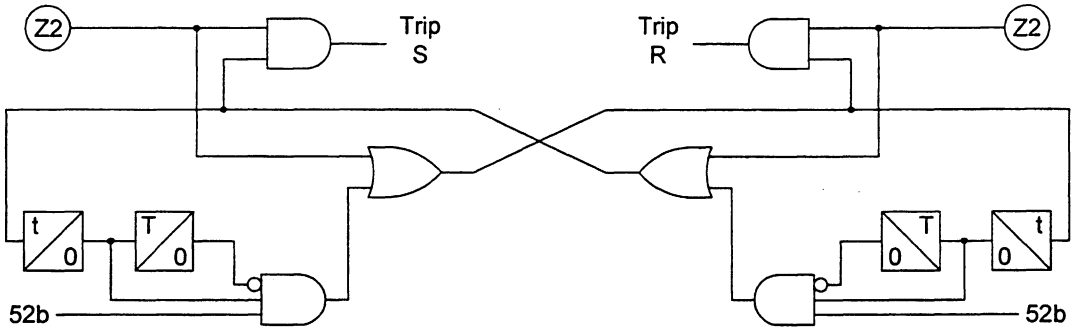

a) Timers and Gate Required to Avoid Channel Lockup

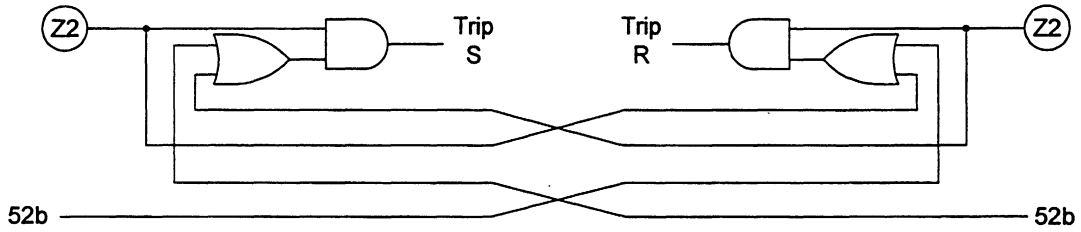

b) Simpler POTT Scheme Eliminates Lockup Problem

Figure 7: Avoid Lockup with Simple POTT Scheme

Control and Balance of Speed, Sensitivity and Security

With multiple bits of information to transfer, we can consider schemes which simultaneously coexist and provide the advantages of each, while minimizing the risk of their individual disadvantages. Figure 8 and Figure 9 illustrate the concept for ground faults. The design uses an individual bit (channel) for each relay element that is desired at the remote end – instead of combining them into one. Then each terminal uses the set of local and remote elements to make its decisions, as desired or required by the immediate operating conditions.

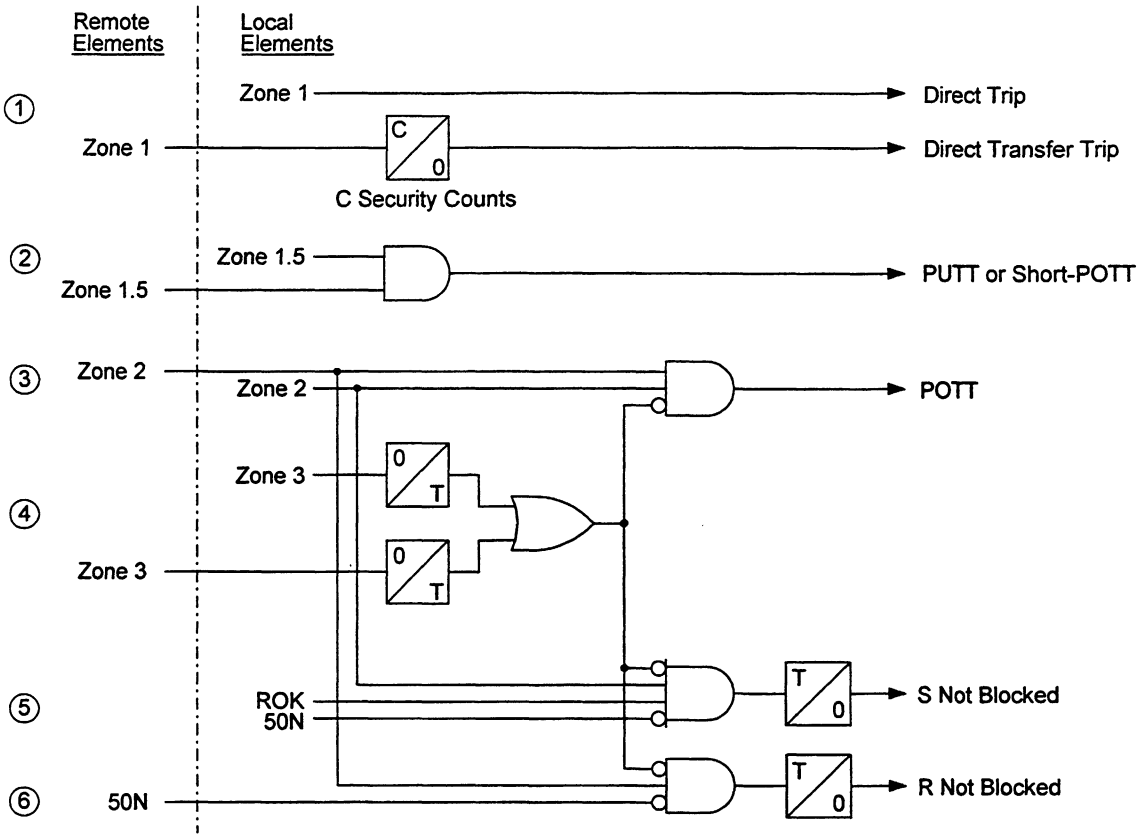

Figure 8: Balance Sensitivity, Speed, and Security

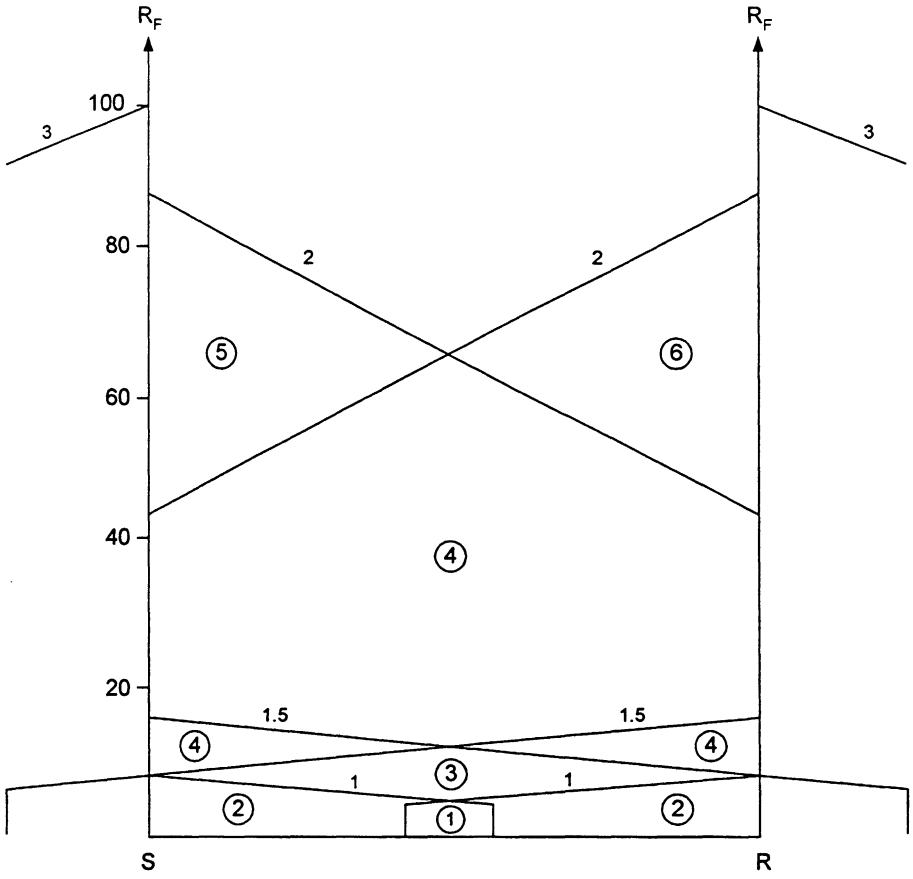

Figure 9: Fault Resistance Regions Covered by the Schemes in Figure 8

The standard IEC 834-1 suggests that the user should decide the states of communicated bits, given a loss of communications. For example, during a loss of communications, should the received bit maintain (hold) its previous state, or default to a 1, or to a 0? As we apply bits in various schemes below, we will also determine the desired action of the bits given communication loss.

① A local Zone 1 element, such as a ground quadrilateral element, trips the breaker directly. Non-pilot tripping covers the area of overlap of the zones 1, marked by the circled 1 in the house-shaped characteristic in Figure 9. We must rely on the channel, or wait for zone timers for all other faults. Appendix A gives some guidelines for setting the resistance and reactance reaches, taking into consideration some of the angle errors that might be expected.

② The remote Zone 1 element state is communicated on one bit and passes through a local security counter to provide a Direct Transfer Trip. The additional regions covered are marked by the circled 2s. Given a channel failure, the received bit should default to a '0'. (A default of '1' would trip the remote breaker for communications failures; a default of 'hold' defeats the receive security counter.)

③ PUTT and short-POTT schemes follow.

In a PUTT scheme (not shown in Figure 8), the local Zone 2 element would be set to reach beyond the remote end. The remote Zone 1 element is an underreaching element. The scheme is fast and simple – and is not encumbered by current-reversal timers or logic, which might delay tripping for faults which evolve to the healthy line.

A short-POTT scheme can also be considered here, and is illustrated in the figures. The Zone 1.5 elements at each terminal would be set to cover the entire line, but would have their reaches limited sufficiently to avoid current-reversals on parallel paths. This improves fault-resistance coverage and speed over PUTT, but does not involve current-reversal timing.

Elements for a short-POTT scheme are labeled 1.5, and cover the additional area marked with the circled 3.

The received bits should default to '0' for communications failures. There is no benefit to holding – only risk.

④ A POTT scheme based on directional overcurrent elements, labeled Zone 2, can provide some more sensitivity, but must be guarded against current reversals. Reverse-looking elements (Zone 3 directional overcurrent elements) and timers provide blocking for a short time after current reverses. The delay that would result if a fault evolved from one line to the other is not important, because of the coverage by the PUTT or short-POTT schemes. A timer could be added after the AND-gate, to gain some security against system unbalances produced by switching, for example. The additional coverage from the POTT scheme is labeled with the circled 4.

Received bits for Zone 2 should default to '0' on loss of channel. A 'hold' or '1' would significantly risk misoperation, with little or no benefit in dependability. The zone 3 received bits should default to a '0,' if communications are lost. A default to a '1' would defeat the protection for T seconds for every bit error.

⑤ A high-resistance fault close to S could be in region 5, which is not covered by the POTT scheme. After some time delay, and if no reverse elements pick up at the remote end, we can trip for such a fault, with the sensitivity we are accustomed to with DCB schemes. However, this scheme enjoys much greater security than DCB.

Conventional blocking schemes assume the fault is internal if the blocking message is not received. (The truth might be that the channel or relay equipment failed to deliver the message to block. Put another way, in traditional DCB schemes, the block signal means "reverse." No block signal means "forward," "not detected," or "bad channel.")

Here, the local terminal knows with near certainty the states of the remote forward and reverse elements, or it knows the channel is down. There is no confusing a lost channel with an internal fault. The fact that we are receiving messages from the remote end reassures us that the remote relay is functioning, and ready to produce blocking signals when appropriate.

In addition, the security can be enhanced somewhat by an undercurrent element, 50N. This element would block tripping, if enough residual current is produced by an open CT, for instance.

The received-OK signal (ROK) ensures this part of the logic is active only if the remote end is successfully communicating, and therefore in a position to block the local end should an external fault occur.

If the source behind R is very weak, then Zone 2 at R might not be able to see all the way back to S. This is the only scheme described here that is capable of detecting the fault.

⑥ A high-resistance fault near R can be cleared in a complementary way, as compared to ⑤ above. The ROK signal is not necessary, assuming the zone-2-received signal defaults to a zero on communications loss.

Protection speeds can be impressive, as the following quick look at a POTT scheme reveals. Consider a direct digital link at 19,200, and allow 6 ms for the message and processing time. Assume the overreaching elements operate in under 12 ms, and the relays use instantaneous trip contacts, such as transistors. Both terminals trip their breakers in less than one cycle. With a 1.5-cycle breaker, the fault is cleared in under 2.5 cycles, at both ends. Given such fast tripping, and also given that the channel could be used for breaker failure, we should investigate shorter breaker-failure times and faster time-step backup.

In summary, the direct tripping Zone 1 elements cover very little of the line when the channel is not available. The coverage is limited to the small region of overlap near the middle of the line. The direct transfer trip path depends on secure measurements from one end, but involves a short security delay of about 4 ms. Because it depends on information from the remote end alone, it can trip the local breaker even if there is a problem at the local terminal such as a loss-of-potential (LOP) condition. The short-POTT path approximately doubles the fault resistance coverage as compared to the DT and DTT paths, but could be disabled by a LOP condition at either end. The POTT scheme adds sensitivity and speed (which might be sacrificed with an extra timer, if temporary unbalances could pick up the sensitive overreaching elements). Again, both ends must determine the fault to be internal. The last two schemes cover faults in the "bow-tie" regions labeled with the circled 5 and 6. Some time delay is required to wait for the possible block from the other end. Overcurrent elements can block these two regions, to ensure they are only active for low-level currents, thereby reducing the risk of misoperation should a current transformer fail at either end.

It should be noted that the schemes above are presented as a concept, to show how communicating multiple bits of information end-to-end can produce an adaptive and balanced protection scheme.

CONCLUSIONS

1. The high quality of many digital communications channels permits more information to be sent in less time, as predicted by Shannon and demonstrated in practice.

2. Direct-digital communications between relays can be designed with the security, speed, dependability, and adaptability needed for blocking, permissive, and direct-tripping applications – as well as for control.

3. Eight bits can be securely communicated every 2 ms, with a worst-case end-to-end delay of 6 ms, including processing latency, over a 19,200 bit/second channel.

4. Extremely secure direct transfer tripping can occur in less than 6 ms, when a two-message sequence is employed over a 19,200 bit/second channel.

5. Direct communications using a simple serial asynchronous bit stream ensures compatibility with a wide variety of communications channels, systems, and test equipment.

6. Channel performance monitoring, including sequence of events, outage duration, and unavailability provides the measurements of performance required to maintain and improve communications – without periodic testing. System operation is a continuous test of the channel.

7. Relay event reports closely relate the performance of the protection and the communications, during faults and other events.

8. Fiber-optic networks and links, digital microwave systems, and point-to-point radio (narrow-band and spread-spectrum) are excellent channels to consider for direct digital-to-digital applications. The combination of lower-cost channels with direct relay-to-relay communications opens up applications at lower voltages, including distribution feeders.

9. Although metallic circuits have demonstrated satisfactory performance in the field, we must consider ground potential rise, isolation, and induced interference during faults. Channel and event monitoring can quickly point out difficulties, should they appear, and lead us to their resolution.

10. High-quality analog channels, such as analog microwave, can be used for digital communications at protection speeds with the help of a limited-delay modem.

11. As always, the channel characteristics, including signal routing, must be considered in selecting and designing protection schemes.

12. Because channels fail in different ways, using different channel types can provide redundancy against failures produced by faults.

13. Communicating eight bits end-to-end opens opportunities for new protection schemes, and for combining traditional schemes for enhanced performance.

14. Schemes can be simpler, and some problems, such as channel lockup, can be avoided or solved more simply when more than one bit is available end-to-end.

15. A POTT scheme implemented over a 19,200 bit/second channel can clear at both ends in less than one cycle, plus breaker times.

REFERENCES

1. Shannon, C.E., 1948, "A Mathematical Theory of Communication," The Bell System Technical Journal, Volume XXVII, July.

2. IEC 834-1, Performance and Testing of Teleprotection Equipment of Power Systems. Part 1: Command Systems, May 1996.

3. OSD 136L, Optical Systems Design Pty Ltd 7/1 Vuko Place, Warriewood 2102, NSW Australia; http://www.optsysdesign.com

4. Haus, Robert J., 1980, *Fiber Optics Communications Design Handbook* pg 158. Inglewood Cliffs, New Jersey: Prentice Hall.

5. Gardner, Terry N., 1998, "Power Networking", Transmission & Distribution World, February.

6. Kahler, J. Kraig, "Installation of 930-960 MHz Low Density Point-To-Point Radios and Solid State Relays for Primary Transmission Relay Protection on 69 kV Transmission Lines," 21st Annual Western Protective Relay Conference, Spokane, Washington, October 18 - 20, 1994.

7. Roberts, J. and Zimmerman, K, "Trip and Restore Distribution Circuits at Transmission Speeds," 25th Annual Western Protective Relay Conference, Spokane, Washington, October 13 - 15, 1998. .

8. AT&T Technical Reference TR 63210, "DS0 Digital Local Channel, Description and Interface Specification," August 1993.

9. IEEE Std 487-1992, IEEE Recommended Practice for the Protection of Wire-Line Communication Facilities Serving Electric Power Stations, November 4, 1992.

10. MBT 9600, Pulsar Technologies, Inc., 4050 NW 121st Ave, Coral Springs, FL 33065; http://www.pulsartech.com.

11. Schweitzer, E.O. III and Roberts, J, "Distance Relay Element Design," 19th Annual Western Protective Relay Conference, Spokane, Washington, October 20 - 22, 1992.

APPENDIX A - QUADRILATERAL REACTIVE REACH VERSUS RESISTIVE REACH SETTING GUIDELINE

QUADRILATERAL ELEMENT REVIEW

To pick up a forward-reaching zone of quadrilateral ground distance protection, the relay must determine that the fault presented to the relay passes the following four measurement test criteria:

- Reactance < set reactance (top line)
- Apparent fault resistance (R_F) < positive-resistance (right-side) blinder
- R_F > negative-resistance (left-side) blinder
- Fault direction is forward as measured by a negative- or zero-sequence polarized ground directional element

Equations 1 and 2 repeat the equations shown in [11] for the Zone 1 A-Phase reactance and resistance tests.

$$X_{AG1} = \frac{\mathrm{Im}\left(V_A \cdot \left(I_R \cdot e^{jT}\right)^*\right)}{\mathrm{Im}\left(Z_{1L} \cdot \left(I_A + k_0 \cdot I_R\right) \cdot \left(I_R \cdot e^{jT}\right)^*\right)} \quad (1)$$

$$R_{AF} = \frac{\mathrm{Im}\left(V_A \cdot \left(Z_{1L} \cdot \left(I_A + k_0 \cdot I_R\right)\right)^*\right)}{\mathrm{Im}\left(\frac{3}{2} \cdot \left(I_{A2} + I_0\right) \cdot \left(Z_{1L} \cdot \left(I_A + k_0 \cdot I_R\right)\right)^*\right)} \quad (2)$$

Where:

Im = Imaginary portion

V_A = A-Phase voltage, [V]

I_A = A-Phase current, [A]

I_R = Residual Current ($I_A + I_B + I_C$), [A]

Z_{1L} = Positive-sequence replica line impedance, [Ω]

* = Complex conjugate operator

Z_{0L} = Zero-sequence replica line impedance, [Ω]

k_0 = $(Z_{0L} - Z_{1L})/(3 \bullet Z_{1L})$, [unitless]

I_{A2} = Neg.-sequence current ($I_A + a^2 \bullet I_B + a \bullet I_C$), [A]

I_0 = Zero-sequence current ($I_R/3$), [A]

T = Non-homogeneous system factor, [degrees]

CALCULATING REACTANCE REACH AS A FUNCTION OF RESISTIVE REACH

The elements described by (1) and (2) are phase angle comparators. For the reactance element described by (1), when the angle between the polarizing quantity (I_R) and the line drop compensated voltage ($Z_{1L} \bullet (I_A + k_0 \bullet I_R) - V$) is 0°, the impedance is on the reactance element boundary. This element must measure line reactance without under- or overreaching from the affects of load flow or fault resistance. Hence, the element must use an appropriate polarizing current: negative- and zero-sequence currents are suitable choices. In some nonhomogeneous

systems, the tip produced by the polarizing current may be insufficient to prevent overreach. To compensate for this nonhomogeneity, we introduce polarizing current angle bias (tip), or reduce the reach of the Zone 1 element.

Reducing the Zone 1 reach restricts that portion of the line protected by overlapping instantaneous Zone 1 protection. This overlapping "zone" is realized for low-resistance faults. As we show next, a large resistive reach can limit the reactance element reach when we consider instrumentation angle errors. If the quadrilateral ground distance elements are the only Zone 1 protection, then we strike a balance between overlapping zone for mid-line faults, and large resistive coverage by one terminal for close-in faults.

Specifically, the instrumentation angle errors we consider are those due to current transformers (CTs), voltage transformers (VTs), and the measuring relay. For our example, the values of these angles are: CT = 1°, VT = 2°, Relay Measurement = 0.2°

Let us consider Relay R shown in Figure A.1.

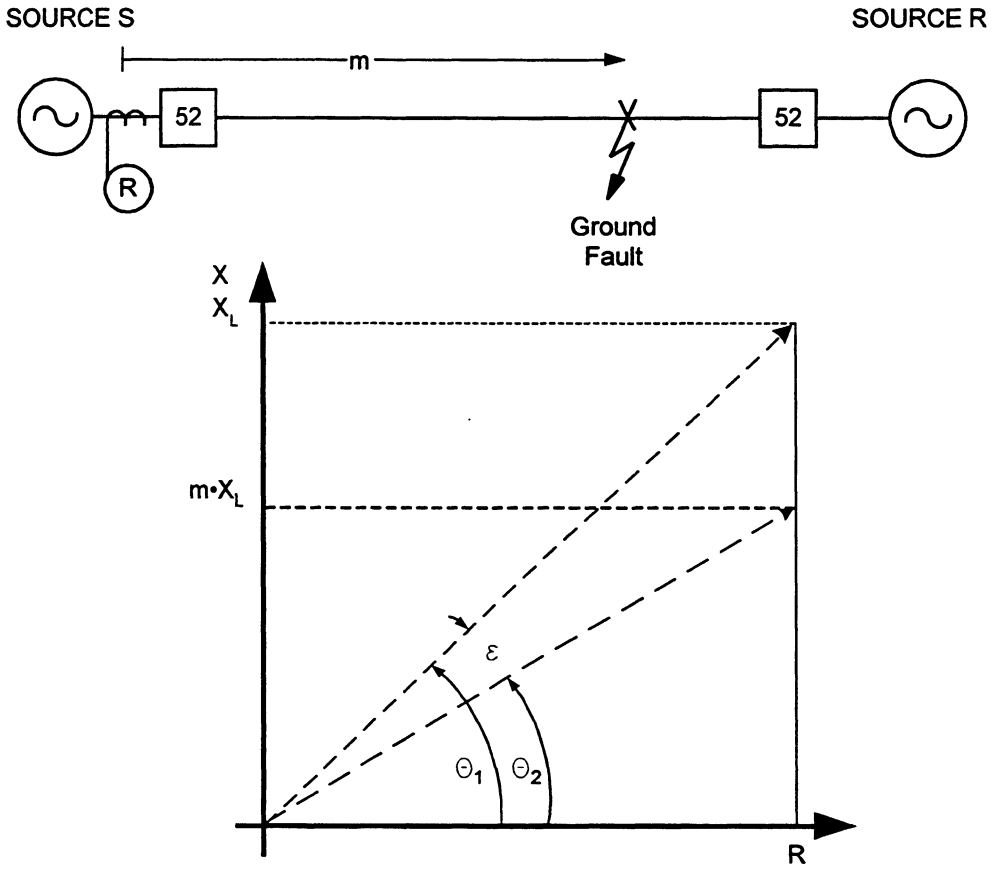

Figure A.1. System-Single Line and First Quadrant of the Quadrilateral Distance Characteristic at Source S Terminal

For a ground fault outside of the protected zone with a reach m [XAG1 of Equation (1)], what is the maximum secure reactive reach for a given resistive reach coverage [RAG1 of Equation (2)]?

From Figure A.1: $\tan\theta_1 = \frac{X_L}{R}$ and $\tan\theta_2 = \frac{m \cdot X_L}{R}$.

Solve for m:

$$\frac{m \cdot X_L}{R} = \tan(\theta_1 - \varepsilon) = \tan\left(\tan^{-1}\left(\frac{X_L}{R}\right) - \varepsilon\right)$$

$$\therefore \quad m = \frac{R}{X_L} \cdot \tan\left(\tan^{-1}\left(\frac{X_L}{R}\right) - \varepsilon\right)$$

For $R >> X_L$, $\tan^{-1}(X_L/R)$ and $\tan(X_L/R) \cong X_L/R$. (Note: this approximation nets an error less than 5% for $X_L/R > 2.5$). If we assume the protected system is homogeneous (i.e, the only angular errors we must account for are those of the CT, VT, and relay), $\varepsilon = 3° \cong 1/20$ radians. Given these simplifications:

$$m = \frac{R}{X_L}\left(\left(\frac{X_L}{R}\right) - \varepsilon\right) = 1 - \frac{R \cdot \varepsilon}{X_L} = 1 - \frac{R}{X_L \cdot 20} \qquad (3)$$

Equation 3 shows us that the lower the resistive reach, the greater the permissible reactance reach. Figure A.2 shows a graph of allowable resistive to reactive reach ratio for $\varepsilon = 1/20$ radians (3°). The dashed line in this figure shows an example where an R/X_L ratio = 8 (for a 1-ohm line and an 8-ohm resistive reach) permits setting m = 0.6 per-unit of the line.

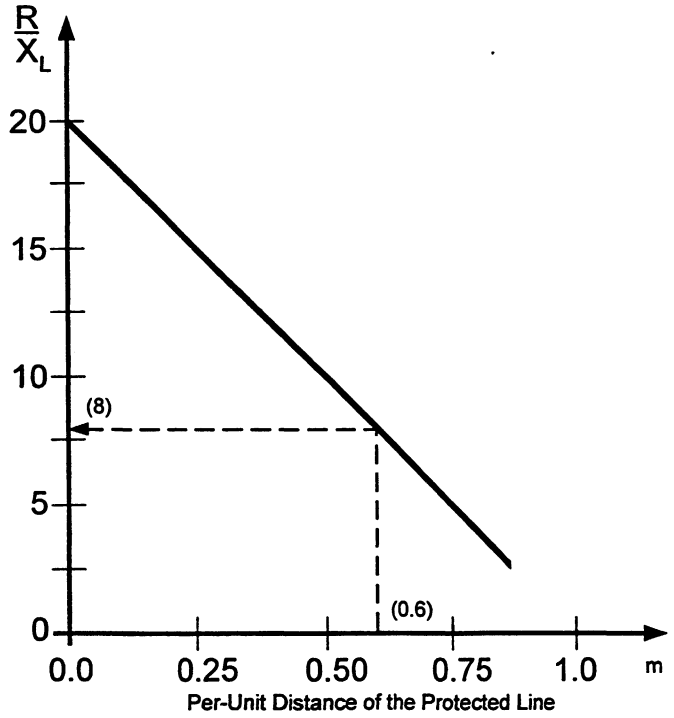

Figure A.2: Increase Reactance Reach By Decreasing Resistive Reach (for $\varepsilon \neq 0$)