

THE SINGLE-PROCESSOR CONCEPT FOR PROTECTION AND CONTROL OF CIRCUIT BREAKERS IN LOW-VOLTAGE SWITCHGEAR

© 2005, IEEE, reprinted with permission from Paper No. PCIC-2003-28, presented at PCIC Houston, 2003 and published in IEEE IAS Transaction 2004

Marcelo E. Valdes, PE
IEEE Member
General Electric Company
41 Woodford Avenue
Plainville, CT 06062
USA

Indrajit Purkayastha
IEEE Member
General Electric Company
41 Woodford Ave
Plainville, CT 06062
USA

Tom Papallo
General Electric Company
41 Woodford Ave
Plainville, CT 06062
USA

Abstract – Protection of low-voltage circuits has been a key area of research and continuous improvement since the invention of electricity. In the current state of the art, fuses and circuit breakers protect individual circuits. These devices, in turn, use many different kinds of trip mechanisms, operating in various modes that require sensing the current that travels through the circuits. Communication networks and external relays are sometimes used to enhance selectivity and provide better protection. In this paper, the authors explore a protection-and-control architecture based on a single processor that provides all the protection and control functions for a lineup of low-voltage switchgear. We demonstrate how this method of protection changes the paradigm from individual circuit protection to system protection, while cost-effectively and significantly increasing system reliability and protection.

Index terms – Low-voltage protection, Single-processor protection architecture.

I. INTRODUCTION TO THE SINGLE-PROCESSOR CONCEPT

A. Background

Communication networks are becoming more common in today's switchgear. They provide the important functions of gathering and reporting information from individual devices, such as trip units, meters, and protective relays. In low-voltage systems, communications often provide supervisory logic, such as load-shedding schemes or reporting of basic electrical information and event status to a central control computer.

A typical architecture with system-wide communications employs a master-slave polling technique. In this type of network, a supervisory computer sends instructions or requests for information to one of the slave devices. The slave device then responds as defined by the network protocol. This implementation has a variable latency, due to such factors as the amount of information requested from each device, the response time of the device, and the communication delay time specified by the protocol. This variable latency limits this type of network to supervisory functions and information gathering. To improve protection, fast, reliable, and deterministic communication is needed.

A varying number of devices and the variable length of each message creates a situation in which response times after

events are difficult to predict and can be relatively slow. Techniques such as interrupts can give a higher priority to information that a sensing device has judged as particularly important. However, this does not lead to a deterministic response with predictable performance. Faster networks, including Ethernet, may provide very fast communication rates, but the collision-detection multiple-access (CSMA/CD) protocols typically employed do not provide predictable, deterministic response times.

Protocols specifically designed for machine and device control, commonly known as Industrial Control Networks or Control Access Networks (CAN), such as Control INET, Profibus, DeviceNET, FIP, Interbus, SDS, ASI, Seriplex, CANOpen, and LON Works were considered. Some of these offered promising capabilities with respect to data rate, reliability, and scalability. However, none offered the optimum combination of capabilities needed to provide fast, reliable, and deterministic communication. Specific derivatives based on the Ethernet standard, such as Ethernet/IP, Profinet, and Modbus/TCP, were also considered by the designers. However, the application layers of these Ethernet derivatives are not compatible with the requirements of the application.

B. The New Concept

The concept discussed in this paper uses a methodology different from that of all other electrical equipment systems to date. Communication is based on the capabilities of Ethernet, while removing the time variation introduced by CSMA/DA protocols. Communication is structured to yield fixed latency and sub-cycle transmission times between a central processor and all the devices in the system. Fast communication and fixed latency are key enablers for using a communication network to perform critical control, monitoring, and protection functions.

A second distinction between this concept and traditional communications found in electrical equipment is in the types of information carried on the network. Rather than processed summary information captured, created, and stored by devices such as trip units, meters, or relays, the actual raw parametric or discrete electrical data and device physical status are carried on the network. The data sent from the devices to the central processor are the actual voltages, currents, and device status. As a result, any microprocessor on the network has complete system-wide data with which to make decisions. It can operate any or all devices on the network based on information derived from as many devices

as the control and protection algorithms require.

The architecture discussed here is centralized, with one microprocessor responsible for all system functions. Alternate architectures, such as distributed and semi-distributed, were considered, but the central processor architecture had the best performance. Fig. 1 shows the centralized architecture applied to a typical lineup of low-voltage switchgear, with a central computer performing the control and protection functions and each breaker acting as a node on the network. The key advantage of this architecture is that the single processor has all the information from all nodes simultaneously. Thus, protection and control schemes can be designed that consider the value of electrical signals, such as current magnitude and phase angle, at one or all circuit breakers in the system with equal ease. This allows the implementation of circuit-specific zone-protection functions as easily as a simple over-current function at a single circuit breaker.

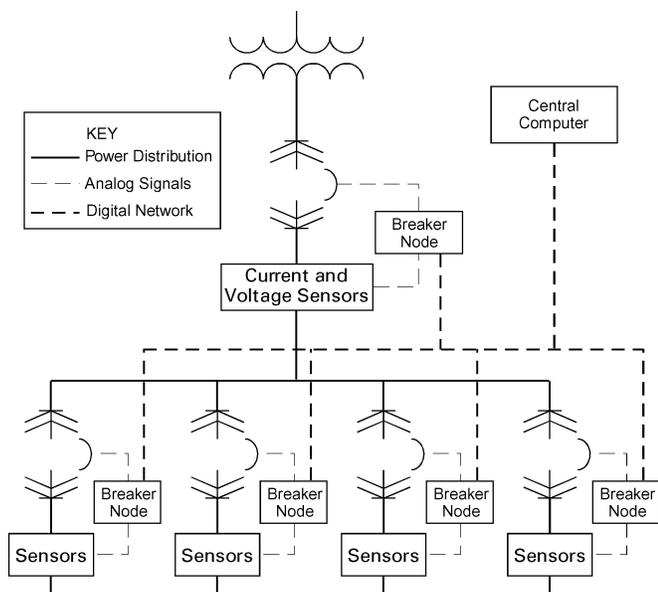


Fig. 1. Centralized control architecture.

This architecture yields fixed latency that includes scheduled communications, a fixed message length, full-duplex operating mode, and a protocol that time slots each message on the network, so that all messages arrive and are handled by the central processor well within a single cycle. This design also schedules all the communication and the programs within the various microprocessors to allow synchronization of information from all devices within a 6- μ s window.

The use of a single processor raises concerns about reliability and a single point of failure that could compromise the protection of all devices within a system. These concerns may be addressed via various methods. The use of hardened industrial electronic components increases the reliability of the individual subsystems. The use of redundant processors can provide increased reliability via redundancy. The electronics at any one circuit breaker node may also be expanded to include back-up protection functions similar to those of a traditional electronic trip. In addition, self-diagnostic algorithms can be employed to detect major or minor subsystem or component

failures and make appropriate adjustments in algorithms to structure the protection into another combination of functions that provide adequate protection. For example, upon loss of a sensor at a main circuit breaker, the over-current functions of that breaker could be directed to function with the sum of the current signals of all the load-side feeder breakers, making appropriate adjustments for the errors that method may introduce. In a traditional system, loss of a sensor at a circuit breaker inhibits all over-current protection at that circuit breaker.

C. Specific Implementation for Low-Voltage Switchgear

Recent advances in processing power and communications bandwidth allow use of such a control system in electrical equipment not only for monitoring, but also for relaying, metering, control, and over-current and short-circuit protection. Each breaker node can digitize the analog signals for each phase at a rate of 32 to 128 samples per 60-Hz cycle, comparable to the rates used in modern trip units. These samples can then be transmitted individually or in groups to the central computer in sub-cycle or even sub-millisecond periods. Network capacities of 100 megabits per second or more enable control of systems of 24 or more breakers by a single computer.

A system of 24 breakers sampling three phase currents, one neutral current, and three voltages at 64 samples per cycle generates data at rates of approximately 10 megabits per second. Not only must the network have the capacity to carry the data, the central computer has to process the data, perform the necessary protection, relaying, and control calculations, then respond to the breaker node before the next transmission. This significant data throughput and processing is possible today using commercially available high-capacity networks and modern microprocessors with computation ability of 1500 MIPS (million instructions per second). Technological advances in the last five years make using this architecture for switchgear protection feasible.

Fig. 2 shows a timeline in cycles for a typical centralized system. The process is started at the node. Each node is scheduled to sample its analog signals at a prescribed rate; 128 samples per cycle in this example. In addition to the data sampling, the individual signals are filtered. Sampling continues at the prescribed period until four samples are collected, at which point the node groups them and transmits this packet to the central computer. The central computer operates on a receive-process-transmit sequence with the same period as the sample transmit. Data from all breaker nodes are received, validated, moved to memory, and preliminary calculations, such as squaring the current, are then performed. Protection-algorithm calculations, as well as metering and other critical processing and logic, are performed. After the calculations are completed, all instructions resulting from the multiple algorithms for the breakers are collected and an instruction message is constructed. This message is then sent to all nodes. The processing sequence is completed before the next transmission from the breaker nodes arrives at the central processor.

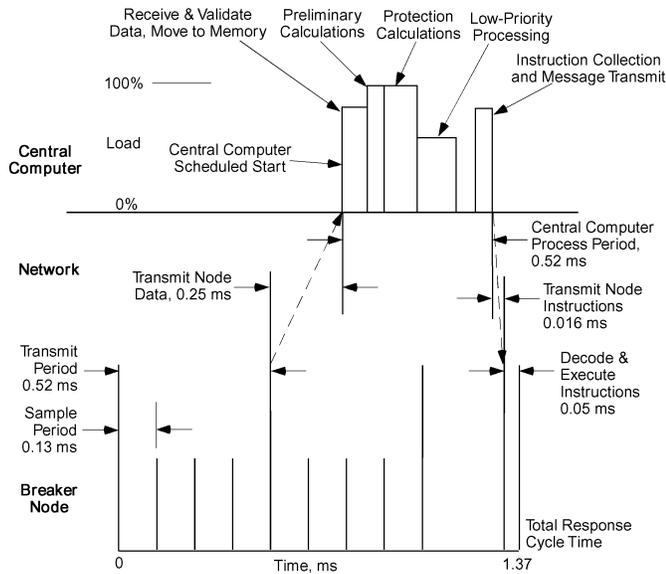


Fig. 2. System processing time line.

Any additional microprocessor capacity can be used to perform lower-priority, non-critical tasks, such as updating kWh and demand logs. The central computer is now ready to receive the next set of data and continue its protection functions. For functions requiring longer data intervals, previous data sample sets are retained and the accumulators are updated with the new data. Using this technique, instantaneous, short-time, and long-time trip functions can be performed well within the times provided by modern integral electronic trips.

When the node receives the information packet, the node decodes it and executes any instructions for changing the circuit breaker state. The timing of this is shown in Fig. 2. The first interval of 0.52 ms is used for data acquisition of four data samples. Data from 24 breaker nodes represents 23,870 bits. Operating at 100 Mbps (megabits per second) the network requires 0.25ms to transmit all node data to the central computer. All algorithms take multiple samples to improve noise rejection. The second 0.52-ms interval is the receive-process-transmit cycle of the central computer. The communication to the breaker node, which is a much smaller message, requires 0.016 ms to reach the nodes. Decoding the message, processing at the node, and activating the analog outputs takes 0.05 ms. Thus, the entire sequence from the initiation of an event to breaker-node response requires 1.37 ms, allowing the sequence to occur 12 times per cycle. The processing cycle time can be accelerated or decelerated for the application's needs by adjusting the sample period and the number of samples per transmission.

This 1.37 ms is the system latency, but does not affect the system synchronization. All of the sample conversions start at the same time. Each message is sequenced to maintain its relationship to the correct time before it is used in a vector calculation. The central computer continuously maintains the data's correct sampling sequence. The order or delays inherent to the calculations do not alter or have an effect on the synchronization of the data to real time.

The advantages of such a system-wide approach may be most appreciated when considering zone-protection functions

that require information from multiple points in the system simultaneously. Each node is simultaneously sampling contemporary data. The central computer can combine the information from various nodes and use it to perform differential-zone functions, such as multiple-source ground fault, zone-selective interlock, and bus-differential protection. Furthermore, in addition to analog electrical information, device status is included in the network data. Knowing the status of all devices, the central computer can determine the system state. Combining the system state with known electrical parameters at each active node creates the possibility of dynamic reconfiguration of the protection functions, according to power flow. This enables dynamic zone protection, with protective settings based on the configuration and loads of the system.

D. Background and Current State of the Art in Circuit Protection

The art and science of coordinating circuit breakers has advanced only incrementally since the first thermal-magnetic breakers and dashpot-operated trip mechanisms were introduced. Bimetallic strips and magnetic forces are still employed in most circuit breakers today. The introduction of electronic trips in the 1960s and of digital trips in the late 1970s have led to the replacement of magnetic-dashpot trip devices and provided an alternative to thermal-magnetic trips. These newer trip mechanisms improved the range and flexibility of settings available with individual circuit breakers. Later developments allowed circuit breakers to signal each other in an attempt to achieve coordinated operation. However, the basic methodology for selectively coordinating circuit breakers has not significantly improved.

Choosing the types of circuit breakers and trip mechanisms and the settings for those trips has always involved trying to find the optimum compromise between system reliability and protection, while operating under investment and space limitations. Better selectivity can be achieved at the cost of lowering the level of protection, while better protection is achieved at the cost of incurring a greater risk of unselective operation or nuisance trips. A good engineer strove to achieve the best compromise satisfying the needs of the system's mission, while keeping within the budget and size constraints.

The most common way to selectively coordinate protective devices is to simply choose devices whose inverse-time characteristics allow the device closest to a fault to detect the fault and initiate a trip, thus interrupting and isolating the fault before a larger line-side device commits to tripping. With thermal-magnetic devices, it is a simple exercise of laying out the trip-time curves so that they do not overlap, also known as nesting the time-current curves.

When molded-case circuit breakers employing magnetic trips are used, the instantaneous or magnetic parts of the curves usually overlap for fault-current values above the highest setting of the line-side device, as shown in Fig. 3. This can result in high-value faults' causing the tripping of multiple breakers simultaneously. This provides protection but clearly sacrifices system selectivity and reliability. Electronic-trip breakers allow more careful shaping of the curves, usually achieving a better compromise between protection and selectivity.

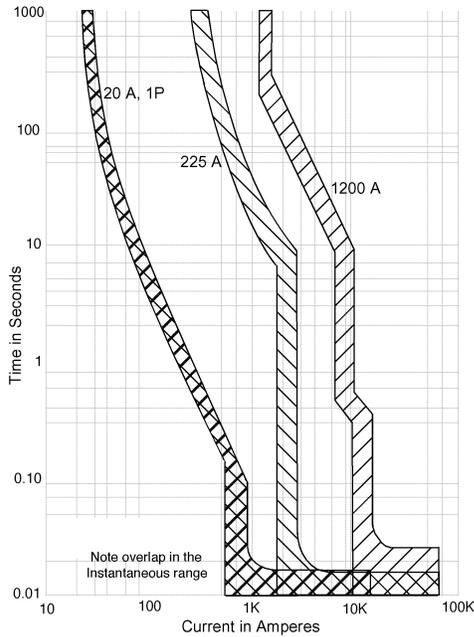


Fig. 3. Time-current curve plot for molded-case circuit breakers.

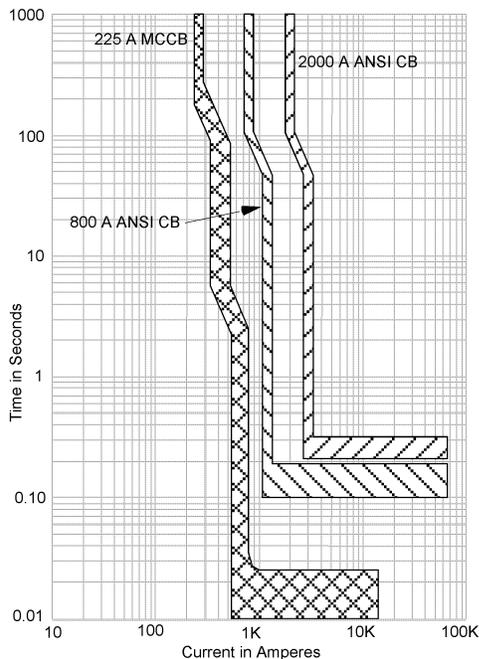


Fig. 4. Time-current curve plot for ANSI circuit breakers.

Newer-design circuit breakers with electronic trips allow the instantaneous settings to be as high as 10 to 15 times the frame or sensor rating of the circuit breaker. However, a 4000-A circuit breaker set at 10X only reaches a nominal setting of 40 000 A, which may be well below short-circuit currents available in a modern power-distribution system. This can result in unselective tripping for high-value faults. Low-voltage power circuit breakers can be used when the lack of selectivity caused by instantaneous protection is undesirable.

The ability of these circuit breakers to delay tripping for up to 30 cycles allows the full coordination of multiple layers of devices within a system, as illustrated in Fig. 4.

Although this achieves much better selectivity, it does so at the cost of sacrificing protection of the conductors and equipment. A further drawback of this method is that a line-side device with several load-side devices under it can only be set as fast as the slowest load-side device and as sensitive as its sensors and zone of protection allow. This means that the line-side device may not be providing the best level of protection for its zone and that it doesn't act as a true backup for all the load-side devices.

E. Ground-Fault Protection

Ground-fault protection presents different challenges and opportunities for circuit protection. A well-understood property of systems in which voltages to ground exceed 150 V is the "arcing ground fault," in which a current conducted through ionized gas can sustain itself indefinitely. The current of an arcing ground fault can be intermittent, initiate a larger phase-to-phase fault, or simply last and cause localized damage. A significant improvement with electronic trips is the ease with which integral ground-fault protection can be provided. Electronic circuit breaker trips typically provide arcing-fault protection by calculating a vector sum of the currents through each conductor in the circuit and comparing the resultant to zero. A value larger than zero implies that current is being improperly conducted to ground. Unfortunately, the inaccuracies of current transformers, charging currents, normal leakage, and the intermittent nature of the currents do not allow this method to be as accurate as one would like. Nevertheless, this method can be used to detect ground currents significantly below the phase currents of the devices. This provides valuable protection from low-current arcing ground faults characteristic of 480/277-V systems. Ground-fault protection systems work well, but often increase the cost and complexity of a protection system.

Various schemes have been used to improve traditional selectivity and protection, most depending on some level of communication among protective devices. The most common scheme used in low-voltage systems is zone interlocking, which is used for ground-fault protection and phase protection in the short-time (100–500 ms) range. If a fault is detected by a load-side device and the device can communicate that fact to the line-side device quickly enough, the line-side device is delayed to allow the load-side device to detect the fault, react, and trip within a set delay time. When the line-side device does not get a blocking signal from any of its load-side devices and it senses a fault, it can react with a faster setting based on the knowledge that the fault is in its zone of protection and not that of one of its load-side devices.

A similar capability is offered in medium-voltage relays, operating within the instantaneous range of the relay. In medium-voltage applications this is usually referred to as "blocking."

A weakness of these methods is that the line-side device does not know which load-side device is sending the blocking signal. Therefore, its fault sensitivity or delay cannot be customized to the specific circuit experiencing the fault. Thus, the line-side device must be set to accommodate the largest fault magnitude and longest time delay of all its load-side

devices. This represents the most significant compromise between selectivity and protection that the system designer must contend with.

F. Differential Protection

Another form of protection involving signals from different points within the system is differential protection. With this scheme, all the currents entering a zone are measured and compared with all the currents properly leaving the zone. When the total currents add to zero, all is well, but if the currents entering are greater than the currents exiting, the difference indicates a fault within the zone, as illustrated in Fig. 5. This method is most commonly applied with bus-differential or transformer-differential relays. This involves the use of separate specialized relays and the proper interconnection of current transformers in every phase around a zone. When multiple zones are used, the locations of the transformers must be carefully thought out so that zones overlap without leaving unprotected sections within the distribution system. The settings of the relays cannot be so sensitive that they are fooled by the multiple inaccuracies and the saturation characteristics of the various transformers. Again, to prevent nuisance trips and overprotection, some level of protection has to be sacrificed. In addition, the cost and complexity of this scheme do not often allow it to be employed in low-voltage power-distribution systems.

II. ADVANCED SYSTEM PROTECTION ENABLED BY THE NEW SINGLE-PROCESSOR CONCEPT

This new system differs from current state-of-the-art low-voltage protection schemes in one critical way. The new system brings all information from all points in the equipment to a single processor with the capability to analyze and control all devices in the system fast enough to support all necessary protection modes, including short-circuit trips of multiple devices simultaneously. The single processor is aware of all important signal values with 6- μ s resolution. This differs from today's systems, which consist of individual protective devices with only minor communication among them. This one distinction significantly broadens the protective capability that the system can achieve economically and reliably.

A. Inverse Time-Current Curves, the Simplest Form of Protection

The simplest form of over-current protection is the inverse time-current characteristics and fixed time delays that can be assigned to each circuit breaker in the system. Even when using these simple inverse time-current time bands to achieve selectivity, the concept system is able to display a significant edge over conventional independent multiple-breaker systems. In a typical main and multiple-feeder system, the main breaker is set to handle the maximum current that the bus may carry, or the sum of the currents of each of the feeders. In the single-processor system, the main would be set likewise with respect to the protection requirements of the main bus. However, the main breaker may also be set with a current setting equal to that of each of the feeder breakers and a time characteristic that allows it to provide backup protection to each individual feeder at that feeder's setting. The processor would simultaneously monitor the current at

the main bus and each of the branch circuits, reacting to an undesirable current at any point. This provides each branch circuit with secondary backup protection optimally set to supplement the primary protection with no compromise needed to achieve selectivity or to allow the bus current to flow unimpeded.

B. Zone Enhancement for the Simplest Form of Protection

In the concept system, the trip characteristics of the breakers can be set as they would in a traditional system, but zone implementations can have a substantial advantage. Traditional zone-selective interlock improves this situation by shortening the main delay, but typically resets the main delay to preset values regardless of which feeder is experiencing the fault.

With the single-processor solution, a zone function replaces the traditional time-current and fixed-delay protection while achieving both selectivity and tight backup protection. The feeder breakers (load-side) are set as before to serve their loads reliably, but the mains and ties (line-side) dynamically set their delay and current settings to best fit each feeder when that feeder circuit experiences a fault.

Consider the example, illustrated in Fig. 5 and Table I, of a double-ended substation, with 4000-A mains, feeding several feeder breakers on either side of a 3200-A tie circuit breaker. For the purposes of this example let us consider a 1600-A feeder on the left bus, Feeder-1, directly fed by Main-1, and an 800-A feeder on the right bus, Feeder-2, fed by Main-2. The Tie is also closed. Typical settings for such a system may have the 800-A feeder breaker set at 2X short-time pickup and a time delay of 60 ms to clear the instantaneous trip of load-side molded-case circuit breakers. The 1600-A feeder could be set at 2X and a time band of 100 ms. Traditionally, the Tie would have a longer delay than the slowest feeder, 200 ms in this case, and the mains would also require a delay slower than that of the tie circuit breaker, 300 ms. With the new concept, the tie and main trip-time curves need not be defined in this manner; the zone function establishes their optimum curves when the fault occurs.

If a fault of 3200 A occurs on Feeder-1, the feeder's protection recognizes this fault and should trip Feeder-1 after a 100-ms delay. Simultaneously, the zone routine recognizes this feeder fault and knows that its location in the system is outside the tie's and mains' zones of protection. The algorithm sets the short-time delay of the Tie and Main-1 to 200 ms and of Main-2 to 300 ms. The Tie and Main-1 will be in short-time pickup mode because of the 3200-A fault current of the feeder, but not for their own currents. When the feeder opens normally, the current drops, the zone function stops timing out, and selectivity is maintained. If the feeder does not clear, the Tie and Main-1 trip at 200 ms to back up the feeder and Main-2 could trip at 300 ms to back up the Tie, if required.

If a fault occurs on Feeder-2, the zone function expects the feeder to clear after a delay of 60 ms. Delays at the Tie and Main-2 are set to 160 ms and Main-1 is set to 260 ms. In this case, selectivity is also achieved. Backup protection is as tight as the circuit breakers allow. In response to the 1600-A fault and the 3200-A fault, the delays of Main-1 and Main-2 are reduced by 180 ms and 220 ms, respectively. It is important to note that, in both these cases, the determination of a fault is based on the feeder's settings and the sensors. In a traditional

system, the sensing of a fault at the tie or main is based on the settings at those trips and the current flowing through the respective circuit breakers. If the current magnitude is not sufficient to be recognized as a fault the trip units will not initiate a trip and hence provide no back-up function.

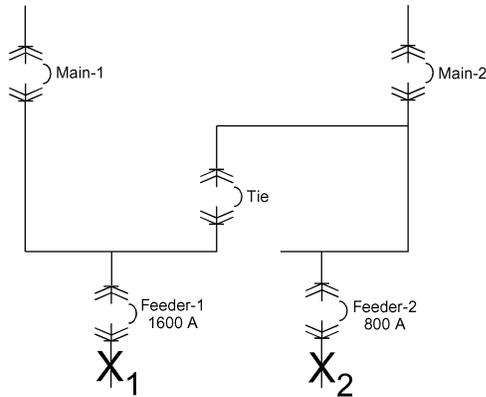


Fig. 5. Example system.

TABLE I.
SETTINGS FOR EXAMPLE SYSTEM

Device	Setting	Time			Current		
		Nominal Setting	Dynamic (Backup Setting 1)	Dynamic (Backup Setting 2)	Nominal Setting	Dynamic (Backup Setting 1)	Dynamic (Backup Setting 2)
Main-1	4000 A	380 ms	200 ms	260 ms	8000 A	3200 A	1600 A
Main-2	4000 A	380 ms	300 ms	160 ms	8000 A	3200 A	1600 A
Tie	3200 A	220 ms	200 ms	160 ms	6400 A	3200 A	1600 A
Fdr-1	1600 A	100 ms	—	—	3200 A	—	—
Fdr-2	800 A	60 ms	—	—	1600 A	—	—

Further, if the fault occurs in the switchgear, the main breaker detects the fault but no feeders experience a fault. The single-processor system is also able to discriminate this situation and can issue a trip command to the proper line-side device, at whatever time delay is programmed, to within one-half cycle. Zone interlocking for the single-processor system is not limited to the short-time delay bands normally associated with circuit breakers. Line-side devices are also aware when a load-side device should be in its instantaneous tripping range and can provide tight backup protection in that situation, as well as for the short-time fault described above.

As can be readily seen, this allows all line-side devices controlled by the single processor to provide perfectly coordinated backup, regardless of the size ratio between the line-side devices and any of the feeder load-side devices. There are no additional margins of safety or unnecessary time delays needed to allow the system to operate selectively and provide protection to the mechanical limits of the devices used.

This scenario also applies within the short-circuit ranges of the devices in the system. When the processor senses a fault within the short-circuit range of any load-side device, the next line-side device is ready to operate immediately if the processor senses that the load-side device is not clearing the fault, even if the fault may not be in the instantaneous range assigned to the line-side device. This form of backup protection could save many cycles of fault current when a

feeder fails to open or if the fault occurs in the switchgear, without sacrificing selectivity.

C. Bus-Differential Protection

The earlier description of bus-differential protection showed it to be a good way to identify faults within the bus of a low-voltage switchgear lineup. However, this type of protection normally requires dedicated current transformers, relays, and wiring, making it too complex and expensive for use on most low-voltage systems. However, this is not the case for the concept single-processor system. Since the processor has the values and vector directions of all the currents at the same time, it is able to measure and detect any current anomaly within the system and can detect an inappropriate lack of current flowing from the feeder. Once the processor has determined that a fault current is originating at the bus, the trip command can be issued to the appropriate circuit breaker immediately. The fault current only lasts long enough to be detected and the time to extinguish is only limited by the mechanical delays of the switching devices. There is no need for the protective algorithms to create any artificial time delays.

Bus-differential protection enhances in-gear fault protection for current levels that traditional protection does not typically cover. Zone-selective interlock reduces the phase-to-phase fault duration for in-gear faults, but only when the magnitude of the fault exceeds the short-time pickup of the main breaker, which can be as high as 36 000 A for a 4000-A breaker. Ground-fault protection operates at much lower currents, as low as 800 A for the same 4000-A breaker, but only provides phase-to-ground protection. Bus-differential protection provides sensitivity in the range of arcing faults for phase-to-phase faults. The only limitation is the combined accuracy of the current transformers or sensors used in the system. The combination of bus-differential protection and zone-selective interlock in conjunction with ground-fault protection presents the opportunity to provide comprehensive arcing-fault protection in switchgear. Ground-fault and bus-differential protection provide phase-to-ground and phase-to-phase protection with sensitive detection levels, able to detect arcing faults. Zone-selective interlocking provides short clearing times at elevated fault magnitudes. All of these techniques, by detecting lower currents or by reducing clearing times, lower the total energy of an arcing fault within the switchgear.

Fig. 6 shows potential timing and current pickups for over-current protection and bus-differential protection on a 2000-A circuit breaker. For a circuit breaker set at a nominal 2000 A, the over-current pickup may start at 2000 A. However, the differential pickup may start as low as 40 A. The current-to-time relationship for the over-current function would follow the typical inverse-time characteristic that includes both long-time and short time delays. However, the differential protection would only need to introduce the minimal delays needed to account for synchronization or sensor timing errors, plus the additional computational time needed to prevent magnitude errors. Use of low-power current transformers with well-known, predictable characteristics can decrease timing errors to well under a millisecond, which can be accounted for within the protective algorithms. The concept system can detect a differential fault in as little as 24 milliseconds.

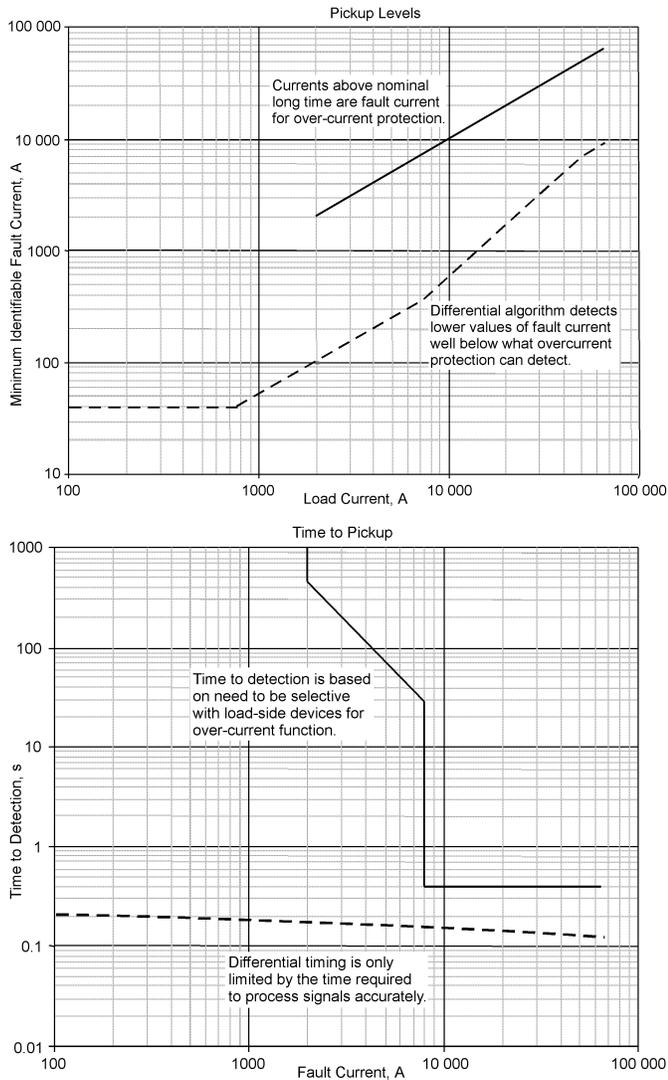


Fig. 6. Differential Protection Timing.

D. Ground-Fault Protection

Arcing ground-fault protection in today's devices also uses time-delay bands and zone interlocking. In the single-processor system, all of the enhancements described for over-current protection are also available for ground-fault protection. Load-side to line-side device curves can be dynamically coupled for immediate backup protection with no unnecessary delay of line-side devices at any detectable fault magnitude. Zone protection allows line-side devices to operate at maximum speed once a fault is detected in the zone. And the differential algorithms can capture a low-current arcing fault whether it is a phase-to-ground or phase-to-phase fault.

E. Simultaneous Multimode Protection

The four modes of protection described so far can be understood as four sets of functions to facilitate understanding of how protection is achieved. However, unlike traditional protective devices, each mode of protection is not provided by

a different device or different sensor. In the concept single-processor system, only one processor is needed for the system's protective functions and only one set of sensors sized for the feeder-circuit current is needed at each circuit breaker.

The single-processor system is able to accommodate more than one processor to provide redundant computing power and is even able to accommodate protective functions at the device nodes for purely redundant backup. However, no protective functions require any additional hardware or wiring beyond the minimum required for any single mode of protection.

F. Monitoring and Event Capture

In state-of-the-art conventional systems, each circuit breaker trip mechanism and meter operates independently. Each device independently captures all data, such as current magnitude, waveforms, or voltages. For multiple devices to capture information simultaneously, an event has to trigger the devices to do so, or a central CPU has to issue a capture command to each device. Typically, available data are limited to current values at trip mechanisms that have opened or indicator flags at those same trips. More advanced trip mechanisms may be able to provide waveform capture with some cycles of information pre- and post-event. Devices that do not sense the event capture no data. When the device closest to the fault does not operate and the line-side device acts as backup, no load-side device information may be available.

A central CPU can issue commands to multiple devices to capture information or the CPU itself can store information from multiple devices when it knows an event has occurred. However, time stamping of such information is not typically done with sub-cycle accuracy. Delays of many milliseconds or even seconds are possible. Resolution can be improved by having the central CPU broadcast synchronization pulses; however, even when that is done, resolution may still be measured in milliseconds.

The single-processor concept system does not share any of these limitations. The time resolution of all information processed is under 6 μ s. The deterministic sequencing of information acquisition by the processor forces information to be continuously synchronized. In a system employing 128 samples per cycle, every sample is synchronized with every other sample in the system.

In addition, during all calculations at the single processor, the information from all the devices is known for the present and for as far back as the algorithm chooses to keep it, at the maximum data-sampling rate. Thus, any event can be programmed to generate the capture and storage of any subset of information desired, including all current magnitudes and phase angles and all setting values and device states. The captured information can extend backward and forward several cycles. The only limitations are those set by the algorithm and the available memory.

Just as protection can be expanded from the limitations of an individual device to a system-wide perspective, event capture and electrical measurement can now be moved from the view available to one device at one point in the system to the expanded panorama seen from an omniscient central point.

III. THE VALUE OF SYSTEM-WIDE INFORMATION CAPTURE, WITH AN EXAMPLE

Consider a system with a double-ended substation designed with 2000-A mains and tie shown in Figure 7. The designer of the system chose not to set up the feeder circuit breakers with ground-fault detection. However, according to code the mains were set up with ground-fault detection and tripping set to nominal values of 0.2X, 400 A, and minimum time delay.

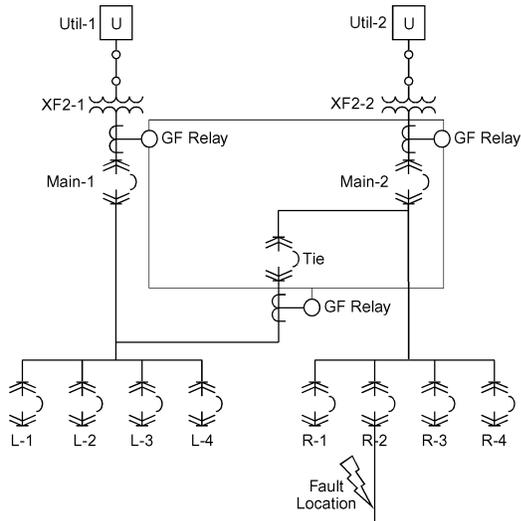


Fig. 7. Double-ended substation example system.

During operation of the system with the tie open, the single-processor measures 2000 A on breaker Main-2 with a slight current imbalance. The single-processor also measures 1600 A with a smaller imbalance on Main-1. The single-processor is capable of performing the ground-fault protection for these breakers because synchronized current from the tie breaker is also available to the processor, capturing the portion of the neutral current caused by the imbalance that is conducted by the tie neutral bus.

During operation, feeder R-2 is conducting 800 A of load current when a 550-A fault occurs between phase A and ground. This fault is immediately detected by the ground-fault relay function, which times to the set delay and trips breaker Main-2. Fig. 8 shows the waveform captured by the relay. If a traditional ground-fault relay with waveform capture were used, this would be the full extent of available diagnostic information. It verifies that a ground fault indeed occurred, but little else.

With the single-processor concept, the single processor is continuously receiving synchronized waveform data from all breaker nodes. Hence, when the ground-fault relay trip occurs, the single processor can store all waveforms for all breaker nodes, even the breaker nodes that did not trip as a result of the fault. Additionally, using a high-speed 32- or 64-bit microprocessor with a large addressable memory space allows several seconds of waveform capture with pre-triggering, rather than just the few cycles available with today's devices.

The value of the system-wide waveform capture for fault diagnostics is immediately obvious by inspection of the

captured waveforms for Main-2 and feeder R-2, Figs. 9 and 10, respectively. Inspection of the Main-2 waveform shows that the Main-2 node detected the incremental 550 A on phase A of the main bus. That is the same 550-A current that can be seen on feeder R-2, on that bus. This allows the system to pinpoint the exact location of the fault current to the phase and feeder location, even though the feeder circuit breaker never tripped. The system operator can now open the suspect feeder breaker and re-energize the main bus to rapidly restore service to the balance of the system.

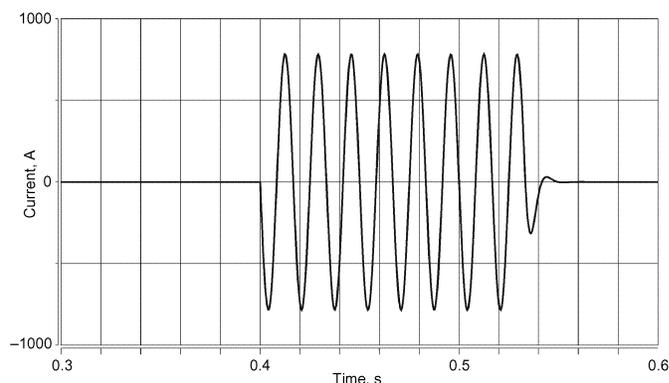


Fig. 8. Waveform captured by example ground-fault relay.

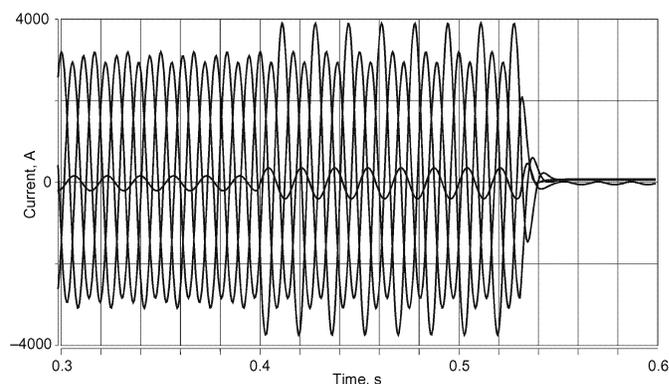


Fig. 9. Captured waveform from Main-2.

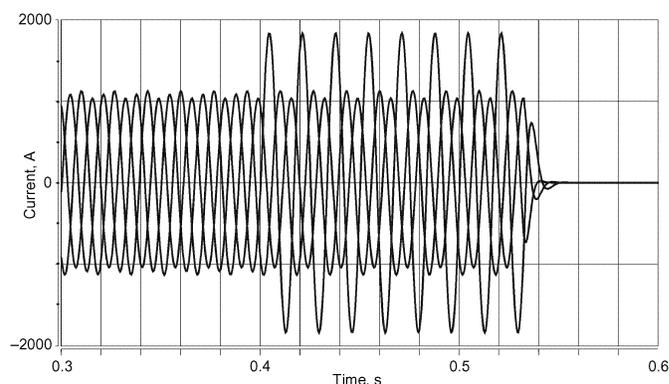


Fig. 10. Captured waveform from feeder R-2.

IV. CONCLUSION

The authors believe that the single-processor concept, coupled with a communication and data-acquisition architecture that allows 6- μ s information resolution, simultaneous capture and processing of all system parameters, deterministic communication and control of all protective devices in the system, and faster than one-half cycle reaction to all fault scenarios will allow significant improvements in electrical power systems protection, reliability and diagnostics.

The benefits can be briefly summarized as follows:

1. Dynamic backup protection of each feeder device by the main device at any current setting without compromising selectivity or protection of either the branch circuit or the main bus.
2. Dynamic zone selectivity that is able to discriminate any current level in the feeder's protection range and adjust time delays at the main breaker to match the needs of that particular circuit regardless of how it is set.
3. Bus-differential protection able to discriminate bus fault currents in the current range of arcing faults with no need for artificial delays to achieve selectivity.
4. Ground-fault protection able to operate at any time delay above one-half cycle without protection compromises to achieve selectivity.
5. System-wide data capture for event diagnostics synchronized as accurately as a single data sample.

Best of all, the single processor is able to cost-effectively achieve all of these functions simultaneously, without encumbering the equipment with multiple protective, metering, and sensing devices, each specialized for a single function.

V. REFERENCES

- [1] Sidnie Feit, *TCP/IP Architectures, Protocols, and Implementation with Ipv6 and IP Security*, New York, NY: McGraw Hill, 1998.
- [2] Stuart R Ball, *Embedded Microprocessor Systems: Real World Design*, Woburn, MA: Newnes, 2000.
- [3] Douglas Comer, *TCP/IP Principles, Protocols, and Architectures*, Upper Saddle River, NJ: Prentice Hall, 2000.
- [4] Bruce Powel Douglass, *Real Time Design Patterns: Robust Scalable Architecture for Real-Time Systems*, Boston, MA: Addison-Wesley, 2003.
- [5] C. Russell Mason, *Art and Science of Protective Relaying*, New York, NY: John Wiley & Sons, 1967.
- [6] J.R. Dunki-Jacobs, "The Effects of Arcing Ground Faults on Low Voltage System Design," *IEEE Transactions on Industry and General Applications*, vol 1A-8, pp 2-8, May-June 1972.

VI. VITA

Marcelo E. Valdes graduated from Cornell University in 1977 with a BS in Electrical Engineering. He has held various positions at General Electric since 1977. He has worked in the GE Installation and Service Engineer organization as a field engineer and service supervisor, and also has also worked in sales, application engineering and marketing for GE Industrial Systems. He currently is the Manager of Application Engineering for the Power Equipment Business in Plainville Connecticut. Mr. Valdes is past chair of the IEEE Power Engineering chapter in San Jose, CA and also the Industrial Applications chapter in San Francisco, CA. He is a registered Professional Electrical Engineer in California.

Indrajit Purkayastha graduated from the Indian Institute of Technology, Kanpur in 1976 with a B.Tech EE degree and from the University of Missouri, Columbia in 1977 with a MSEE degree. He has held various design and management responsibilities in Technology with different businesses within General Electric Company since 1978. He is currently responsible for the Power Equipment Business Technology, Plainville, CT. He is an author of two previous PCIC papers.

Tom Papallo graduated from the University of Connecticut in 1986 with a BS degree and in 1989 with an MS degree, both in Mechanical Engineering. He started with GE in 1986 in the New Product Development department. He has also worked on circuit breaker and electrical distribution system projects for several other major manufacturers, returning to GE in 1997. He is currently the technical lead for a New Product Development project and an adjunct member of the Design Office for the New Product Introduction Department, Plainville, CT. He is the holder of 18 US and international patents.