Best of Both Worlds – Analog Principles Applied in a Digital Distance Relay

Bogdan Kasztenny, Mangapathirao V. Mynam, Chad Daniels, Titiksha Joshi, and Amol Kathe
Schweitzer Engineering Laboratories, Inc.

Revised edition released June 2021

Previously presented at the
47th Annual Western Protective Relay Conference, October 2020

Previous revised edition released September 2020

Original edition released May 2020

Best of Both Worlds – Analog Principles Applied in a Digital Distance Relay

Bogdan Kasztenny, Mangapathirao V. Mynam, Chad Daniels, Titiksha Joshi, and Amol Kathe,
Schweitzer Engineering Laboratories, Inc.

***Abstract*—Early microprocessor-based distance relays did not abandon the state-of-the-art time-domain coincidence timing methods available in the 1980s because of their substandard performance, but because the limited processing power of the early microprocessor-based distance relays did not allow them to use the analog methods. Microprocessor-based relays, digital protection, and phasor-based operation became synonymous. It is time to revisit this notion. This paper reviews the basics of coincidence timing for shaping distance element characteristics (mho and quadrilateral), explains the benefits of using coincidence timing, shows digital implementations that far surpass the dreams of analog relay designers, and presents test results from an implementation that uses the best of both worlds – analog principles implemented in a digital relay.**

Preface

This paper comprises material taken directly from two short papers that we recently wrote to address two individual aspects of a more general theme.

Sections II, III, and IX through XI are directly based on [1] and explain the basics of coincidence timing and the many advantages of this approach. The material explains why early digital relays could not afford this implementation method and why they settled on using heavily filtered phasors. The material focuses on describing a modern digital implementation of distance protection elements based on coincidence timing that combines the advantages of the core analog principle with several improvements that take advantage of today's digital relay technology.

Sections IV through VIII are directly based on [2] and present the theory, implementation, and laboratory test results of a new filtering method for protective relaying based on window resizing. The method uses a full-cycle sliding data window until a disturbance is detected. Upon a disturbance, the window size is considerably shortened to include only disturbance samples and exclude all predisturbance samples. Over time, the window size grows to include more disturbance samples as they become available. When the window reaches its nominal full-cycle size, it stops extending and starts sliding again. By purging the predisturbance data, the new filter strikes an excellent balance between speed and accuracy. We derive the method for fixed sampling and processing rates and include compensation for frequency deviations in the input signals.

I. Introduction

Recently, new implementations of protection elements and schemes became available that are based on superimposed components and traveling waves [3] and operate as fast as 2 ms for line faults. These protection principles work on fault-induced signals supplied not only by the sources but also by the network, by using the energy stored in the LC network parameters prior to the fault. This reduced dependence on power sources makes these new protection principles a viable solution for protection applications near nontraditional power sources. However, transient-based protection methods are not fully dependable because traveling waves dissipate and incremental quantities expire. Therefore, transient-based protection elements and schemes need a dependable backup.

Dependable protection elements and schemes must work on the same signal spectrum as the power sources that drive fault currents in the grid so that the sheer existence of the fault provides lasting operating signals for these principles. Modern power grids with high penetration of wind generators, inverters, and static condensers supply fault currents only for a short time. Therefore, backup for transient-based protection must not only be dependable, but also fast. Otherwise, the power-frequency signals may also be squelched by the nontraditional sources and seriously challenge protection dependability. Considering nontraditional power sources, protection speed is an important way to improve protection dependability.

Reduced inertia of present power systems makes it more difficult to ensure transient power system stability, and it demands faster relay operation. Public safety and wildfire hazards are also important factors in the renewed interest in faster protection operating times.

Historically, protective relays use band-pass filters to obtain protection operating signals consistent with the frequency of fault currents and voltages, while rejecting other signal components (phasor-based protection). Since the beginning of protective relaying, relay filter designers have strived to address the contradicting requirements of speed (short group delay of the filters) and security (accurate measurement through rejection of the out-of-band signal components). The electromechanical relay technology limited the designer options for speed because of the inherent inertia of an electromechanical apparatus at the heart of a relay system. This inherent inertia added a degree of delay equivalent to extra filtering.

Unit protection principles (differential and directional comparison schemes) are inherently fast and can operate as fast as half a power cycle. Speeding up the directly tripping distance and high-set overcurrent elements is considerably more difficult. Diminishing short-circuit current levels will prevent us from applying high-set overcurrent elements. As a result, our focus shifts to distance elements for line protection, especially

the directly tripping Zone 1 elements that operate without a protection channel.

In the 1970s, when semiconductor components became reliable enough for protective relay applications, relay designers introduced distance relay designs based on filtering and coincidence timing. These designs used analog circuits with semiconductors. Free of the inherent inertia of electromechanical devices, these distance relays operated very fast. As the industry continued to learn about electromagnetic interference and semiconductor failure modes, microprocessor-based relays disrupted the field with unprecedented flexibility, new functionality, and unparalleled self-monitoring to mitigate hardware failures. Manufacturers and users moved on to the digital technology, and static relays became a "lost generation."

The early digital relays could apply only very limited sampling and processing rates. Out of necessity, these relays abandoned the time-domain approach of static relays and started a new path for implementing protection functions. This new path focused on "slowing down" the flow of information so that early microprocessors could keep up. These relays applied heavy low-pass filtering in order to be able to sample just several times a cycle. They "compressed samples into phasors" at the front end of the processing chain for the key benefit of processing phasors at relatively low rates. Even today, many microprocessor-based relays process protection logic just a few (four or eight) times a cycle.

Early microprocessor-based relays did not abandon time-domain coincidence timing because of its substandard performance, but rather because their limited processing power did not allow them to use analog methods. Over the first three decades of microprocessor-based protective relaying, digital protection and phasor-based operation became synonymous. It is time to revisit this notion. New digital relays have enormous processing capabilities. High sampling and processing rates now allow us to implement and improve principles invented for static relays.

Any comparator in a distance protection element can be shaped by comparing the angles (coincidence) of the operating and polarizing signals. For example, a positive-sequence-polarized mho comparator uses $S_{OP} = I \bullet Z_R - V$ and $S_{POL} = V_1$, where I and V are the measured current and voltage, Z_R is the reach impedance, and V_1 is the positive-sequence voltage. These two signals can be compared in either the frequency domain (phasors) or the time domain (coincidence timing). Regardless of the implementation (heavy filtering with phasors or moderate filtering with low-pass filters), protective relays always apply some degree of filtering and use the filters to control the balance between speed and security.

Today, microprocessor-based protective relays use finite impulse response (FIR) filters with sliding data windows for band-pass filtering and measurement of phasors. Cosine or Fourier data windows are commonly used. Short windows yield faster operation but allow larger transient errors. Often, intentional delay, reduced reach, and additional restraining in general, are used in the protection logic to address these transient errors. These methods, however, partially or entirely erase the initial gain of speed and make the design less effective. In some cases, two parallel measurement paths are used, such as with full-cycle and half-cycle (or even quarter-cycle) filters operating in parallel. The full-cycle measurement is slower but dependable. The measurement with a shorter data window is faster, but it may be intentionally desensitized and may operate only under certain favorable conditions.

This paper presents the theory, implementation, laboratory test results, and a field case example of a new filtering method that uses filter window resizing to achieve the following:

- Speed of operation.
- Accuracy of the operating characteristics.
- Efficiency of implementation.

Highlights of the new method include the following:

- The filter window resizes to a short length upon detecting a disturbance. The window subsequently grows with each new available sample and eventually slides after reaching its full nominal length.
- The filter includes a carefully designed resizing logic to allow or prevent resizing in order to provide optimum performance while maintaining security.
- Instead of letting the window slide from the pre-fault state to the fault state, the algorithm intentionally delays window resizing for a few milliseconds so that the predisturbance data are entirely purged from the filter window. Hence, the shortened window contains only the fault-state data. This method provides good accuracy, despite using a short data window following resizing.
- The method can be used as a plain filter, and therefore it may be combined with any other post-processing algorithm, such as a pair of orthogonal filters for phasor estimation. Or, the method can be used directly to measure phasors.
- The method is applicable to protection elements and schemes that use fundamental frequency measurements across all protection applications.
- The method is derived for relay hardware with a fixed sampling rate and an arbitrary fixed processing rate.
- The method compensates for off-nominal frequency of the inputs and for the group delay, allowing the downstream protection logic to compare the input samples with the filter output samples, if desired.

After presenting this novel filtering method, the paper presents a new distance element design based on coincidence timing, explains the benefits of using coincidence timing, and shares some key details of digital implementation and improvements that far surpass the dreams of analog relay designers. Finally, it presents test results for a distance relay implementation that uses the best of both worlds: analog principles implemented in a microprocessor-based relay.

II. Distance Protection Element Overview

In general, a distance protection element consists of several logical conditions (comparators) joined with an AND gate. For example, a quadrilateral distance element includes a reactance comparator, a right blinder comparator, a left blinder

comparator (optional), a directional comparator, and a faulted-loop selection comparator. A mho distance element includes a mho comparator, a faulted-loop selection comparator, and a directional comparator. The mho distance element can be further modified by optionally adding a reactance comparator or a blinder comparator.

The performance of all individual comparators that make up a distance element is important for the overall performance of the element. However, the speed and security of a distance element is mostly affected by what we refer to in this paper as "reach-sensitive comparators." Reach-sensitive comparators are responsible for distinguishing between faults located short of the reach point (element operates) and faults located beyond the reach point (element restrains). All other comparators assert for a fault short of and beyond the reach point, and it is only the reach-sensitive comparator that decides if the element operates or restrains.

These reach-sensitive comparators are the mho comparator in the mho distance element and the reactance comparator in the quadrilateral distance element. To some degree, the blinder comparator (resistive reach comparator) in the quadrilateral distance element is also a reach-sensitive comparator.

Historically, a distance comparator is explained by using two signals: an operating signal (S_{OP}) and a polarizing signal (S_{POL}). When enough filtering is applied, the two signals are sine waves (during faults or in a no-fault state). A comparator asserts its output if the S_{OP} and S_{POL} signals are approximately in-phase, and it deasserts if the S_{OP} and S_{POL} signals are approximately out-of-phase. Typically, the operating threshold is drawn at 90 degrees. If the angle between the S_{OP} and S_{POL} is between –90 and 90 degrees, then the comparator asserts. If the angle is outside this interval (–90 to 90 degrees), the comparator deasserts. For example, a mho comparator uses the following signals:

$$S_{OP} = I \bullet Z_R - V \tag{1a}$$

$$S_{POL} = V_{POL} \tag{1b}$$

where:

- I is the relay current.
- V is the relay voltage.
- Z_R is the reach impedance (setting).
- $I \bullet Z_R$ represents a voltage drop across the intended reach impedance Z_R.
- V_{POL} is the polarizing signal, such as the relay voltage (self-polarized mho), healthy phase voltage (cross-phase-polarized mho), positive-sequence voltage (positive-sequence-polarized mho), or a pre-fault voltage (memory-polarized mho).

The V and I terms are adequately selected from the three-phase quantities (V_A, V_B, V_C; and I_A, I_B, I_C) based on the fault type. Today, six instances of the comparator are implemented to monitor all six protection loops for the AG, BG, CG, AB or ABG, BC or BCG, and CA or CAG faults. For any given fault type, the distance element permits only some loops to operate. A faulted-loop selection comparator is responsible for deciding which loops are permitted to operate.

The S_{OP} and S_{POL} signals can be developed in the time domain or frequency domain, as follows:

- *Time-domain* implementation uses the R•i + L•di/dt term to replicate an instantaneous voltage drop across the reach resistance and inductance, subtracts it from the instantaneous voltage, and obtains an instantaneous operating signal (S_{OP_INST}).
- *Frequency-domain* implementation uses the I•Z term to calculate a voltage drop phasor across the reach impedance, subtracts it from the voltage phasor, and obtains the operating signal phasor ($\overline{S}_{OP}$). Alternatively, a frequency-domain implementation can pass the instantaneous operating signal (S_{OP_INST}) through a phasor estimator to obtain the $\overline{S}_{OP}$ signal.

The reactance comparator uses these signals:

$$S_{OP} = I \bullet Z_R - V \tag{2a}$$

$$S_{POL} = j \bullet I_{POL} \tag{2b}$$

where j relates to a phase shift by 90 degrees in the frequency domain or the di/dt operation in the time domain.

We can obtain various reactance comparators by using different polarizing currents. For example, I_{POL} can be the loop current I (self-polarized reactance), the negative-sequence current I_2 (negative-sequence polarized reactance), or the zero-sequence current I_0 (zero-sequence polarized reactance). Note that the mho and reactance comparators only differ by the polarizing signal they use, while their operating signals are identical.

Various relay technologies check the angle between the S_{OP} and S_{POL} signals differently:

- *Electromechanical relays* are designed to develop a torque from the S_{OP_INST} and S_{POL_INST} signals to move the relay rotor in the operating direction if the torque is positive (i.e., the angle between the two signals is between –90 and 90 degrees) and close a contact.
- *Microprocessor-based relays* that use phasors can follow one of the following three approaches (Fig. 1):
  a) Calculate the angle directly and check it against the 90-degree threshold.
  b) Calculate the torque and determine if it is positive $Re[S_{OP} \bullet conj(S_{POL})] > 0$.
  c) Calculate the fault distance m-value and determine if it is lower than the reach impedance setting.

All these phasor-based implementation methods are mathematically identical. They only differ in terms of required operations and computational burden. For example, the m-value method is computationally very efficient when implementing multiple zones with identical settings, except the reach setting.

Static relays use coincidence timers to check how long the S_{OP_INST} and S_{POL_INST} signals are of the same polarity. After low-pass filtering, the S_{OP_INST} and S_{POL_INST} signals are sine waves. If they are perfectly in-phase, they coincide (have the

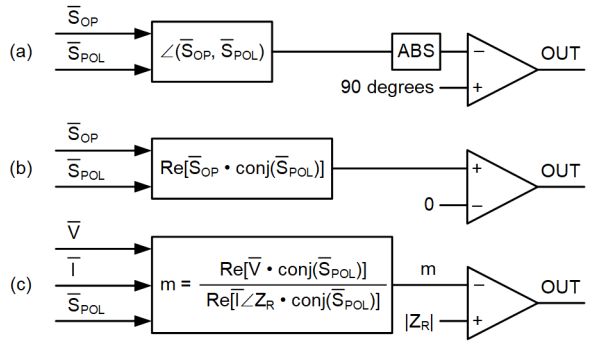

Fig. 1. Implementations of a distance comparator in phasor-based microprocessor-based relays.

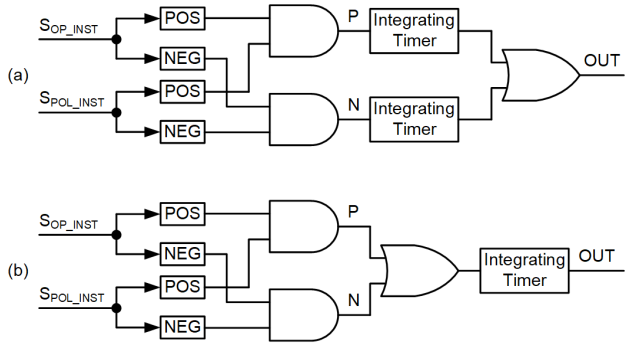

Fig. 2. Implementations of a distance comparator in static relays: (a) dual-timer method and (b) single-timer method.

same polarity) for a half cycle in each half cycle. If they are 90 degrees apart, they coincide for a quarter cycle in each half cycle. If they are perfectly out-of-phase, they do not coincide at all. Static relays use rectifier circuits to detect the instantaneous polarity (sign) of the S_{OP_INST} and S_{POL_INST} signals, a few AND and OR gates to detect if the S_{OP_INST} and S_{POL_INST} signals are the same polarity, and a timer to determine if the matching polarity intervals last for longer than a quarter cycle (90-degree coincidence), see Fig. 2.

The coincidence timing method has several advantages:

- Transients that may be present in the operating and polarizing signals can be used for security. For example, if during the matching polarity period a moment of opposite polarity occurs, the timer can reset or integrate down, providing extra security.
- The comparator operation is fast; it takes only a quarter cycle to detect the 90-degree coincidence.
- An independently designed and optimized filtering scheme can be applied to voltages and currents before passing these signals for coincidence timing.
- Some transients that occur in voltages and currents have a chance to mutually cancel in the $S_{OP} = I•Z – V$ signal. The scheme does not need to excessively suppress transients in voltages and currents separately, but it can focus on transients in the operating signal.
- The $S_{OP} = I•Z – V$ signal can be inspected for level: large signals indicate faults away from the reach point (internal or external); small signals indicate faults close to the reach point (internal or external). Adaptive levels of security can be applied based on the S_{OP} magnitude.

The coincidence timing method can also bring additional benefits to speed and security. Fig. 3 illustrates operation of three different versions of an integrating timer. The design of Fig. 3a is biased toward security: any momentary dropout of the input restarts the timer. Fig. 3b is biased toward dependability: a momentary dropout results in holding the integrator (for a finite time), and when the input picks up again, the integration starts from where it stopped. Fig. 3c is a hybrid solution: when the input deasserts, the integrator does not reset instantaneously but integrates down, away from the operate threshold and toward a complete reset. A modern digital implementation can

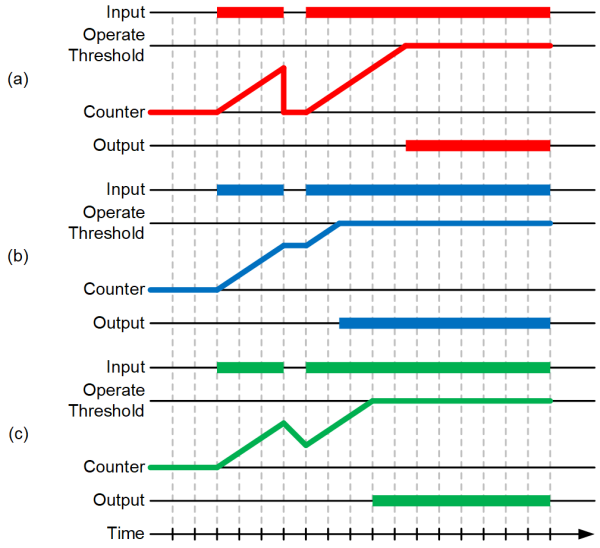

Fig. 3. Integrating timer: (a) instantaneous reset, (b) hold, and (c) integrate down.

use any of these solutions or switch dynamically among them, depending on other conditions.

Early microprocessor-based distance relays could not afford fast sampling and processing. They were incapable of emulating the coincidence timing method. Instead, they used phasors. Some relays further distilled phasors into either the angle between the operating and polarizing signals or into the m-value. These methods lose the association with the I•Z – V signal and the related benefits listed previously in this section.

III. Signal Processing and Filtering

Our new filtering method derives instantaneous operating and polarizing signals in the time domain following the first principle of distance protection for a three-phase power line. The method uses a current derivative to obtain the I•Z terms (instantaneous voltage drops across the line replica impedance), which is a classical solution dating back almost a century to the first electromechanical distance relays. Reference [3] provides details on the numerical implementation of the line replica circuit.

We low-pass filter the instantaneous operating and polarizing signals with a second-order infinite impulse response (IIR) filter to reject high-frequency signal components that

would otherwise violate the RL line model we used in the numerical line replica circuit. The –20 dB point of this low-pass filter is set at a frequency of several hundred hertz.

Finally, we apply a variable-window FIR filter to obtain the direct and quadrature components of the instantaneous operating and polarizing signals. This variable-window filter with window resizing is a key contributor to the element operating time. Think of this filter as a nonstationary, fast, and accurate phasor estimator, outputting the real (direct component) and imaginary (quadrature component) parts of the input signal.

IV. Filter Window Resizing

Fig. 4 illustrates the concept of window resizing. The input signal transitions from a predisturbance state (blue trace) such as a load current, to a disturbance state (red trace) such as a fault current. The method applies a disturbance detection logic to identify the presence and time of the disturbance. A variety of approaches can be used for disturbance detection. Typically, a change-over-time approach is used, such as by comparing input signal samples over one cycle of the fundamental frequency, or by comparing zero- and negative-sequence phasors over an arbitrary time interval, such as one cycle or a half cycle. When using relatively high sampling rates, on the order of a few kilohertz, we can reliably detect the presence and time of the disturbance with submillisecond accuracy.

Our method uses a window-resizing logic to initiate window resizing following the disturbance. The resizing logic includes several security conditions to prevent resizing under unfavorable conditions, such as when the relay does not measure system frequency or during a breaker operation for an external fault. The resizing logic initiates resizing from the full window length of N samples (such as one cycle), to the minimum window length of M_0 samples (such as one-tenth of a cycle). The logic delays resizing to make sure that the initial window excludes the predisturbance samples and only includes samples associated with the disturbance state. The intentional delay between the DISTURBANCE bit and the window RESIZE bit may account for the group delay of low-pass filters in the relay, as well as the inherent low-pass filtering behavior of instrument transformers, especially capacitively coupled voltage transformers (CCVTs).

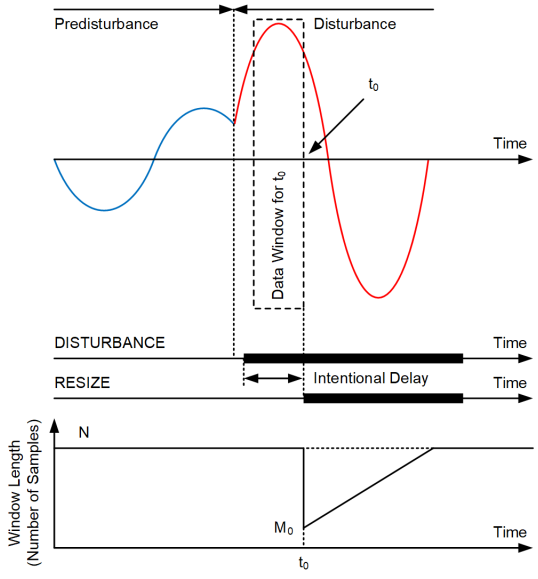

Fig. 4. Illustration of window resizing.

When the RESIZE bit asserts at time t_0, the short data window includes only disturbance samples. The oldest sample in the window may be about 2 ms past the predisturbance samples to avoid polluting the window with a transition from the predisturbance state to the disturbance state.

Our analysis and simulations show that we obtain better protection operating times by using a short data window and delaying resizing than when using a longer data window and allowing it to slide across the predisturbance-to-disturbance transition time. For example, a 2 ms data window starting 4 ms into the disturbance (data between 4 and 6 ms into the fault) gives better results than a sliding data window of 4 ms (data between 0 and 4 ms into the fault).

V. Filter Block Diagram

Fig. 5 shows the block diagram of the variable-window filter. The filter works with the input signal x and outputs the filtered signal y. The filter logic compensates the output y for the variable group delay; therefore the x and y signals are time-coherent and can be directly compared if the downstream protection logic requires it. Alternatively, the filter can be implemented as a phasor measurement algorithm with the input x and a complex (phasor) output x_C.

The *Resizing Logic* subsystem controls the window-resizing operation and provides the filter logic with the present window length M. After resizing takes place, M is a small fraction of a one-cycle window, and it grows to a fixed window length of one cycle. At this point, the window stops growing and starts sliding.

The filter uses a pair of orthogonal FIR filters with the direct filter having window coefficients denoted as h_D and the quadrature filter having window coefficients denoted as h_Q. This paper uses cosine- and sine-shaped windows for the h_D and h_Q coefficients, respectively. Our method, however, can be extended to any pair of orthogonal filters.

The *Frequency Measurement* subsystem provides the filter logic with the present value of frequency f. The filter logic requires the system frequency to follow the frequency of the input signal. The variable-window filter is a band-pass filter, so

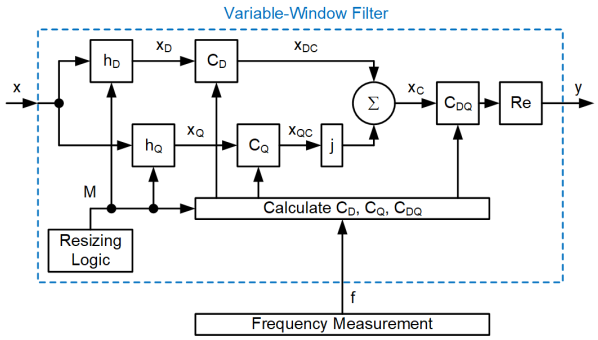

Fig. 5. Block diagram of the variable-window filter.

it must follow the frequency of the input signal. We implemented and tested the variable-window filter with the frequency measurement subsystem described in [4].

The orthogonal components of the input signal, x_D and x_Q, obtained by using the h_D and h_Q filter windows, respectively, require gain-compensation accounting for both the present window length M and the present system frequency f. The C_D and C_Q multipliers provide the required gain compensation.

After resizing, the left edge of the data window is fixed, while the right edge of the data window advances with each new available sample. As a result, the midpoint of the data window advances through half of the sampling period with each new sample of the input signal x. When the window reaches its full length and starts sliding, the midpoint of the data window advances one sampling period with each new sample of the input signal x. This behavior creates a variable group delay. The filter uses a complex multiplier (phase shifter) C_{DQ} to shift the x_C phasor in such a way that the group delay between the input x and the output y is eliminated. The phase-shifting multiplier C_{DQ} is also a function of both the present window length M and the present system frequency f.

In general, the correcting coefficients C_D, C_Q, and C_{DQ} are complex functions of two variables: M and f. This paper derives simple, yet adequately accurate, approximations of these correcting coefficients for the cosine- and sine-shaped orthogonal filters. The paper also teaches how to derive these coefficients for any pair of orthogonal filters.

VI. Filter Implementation

A. Coefficients of the Orthogonal Filters

We use the following cosine- and sine-shaped windows for the orthogonal filters h_D and h_Q:

$$h_{D(k)} = \cos\left(\frac{2\pi\cdot(k-0.5)}{N}\right), k = 1..\frac{N}{2} \tag{3a}$$

$$h_{Q(k)} = \sin\left(\frac{2\pi\cdot(k-0.5)}{N}\right), k = 1..\frac{N}{2} \tag{3b}$$

where N is the window length, calculated as follows:

$$N = 2\cdot\left(\text{floor}\left(\frac{0.5\cdot f_S}{f_{NOM}}\right)+1\right) \tag{4}$$

where:

f_S is the relay sampling frequency (Hz).

f_{NOM} is the system nominal frequency (Hz).

We assume the relay samples and processes data at fixed time intervals (f_S is constant and not a function of power system frequency). We prefer fixed sampling and processing rates because they simplify the hardware and firmware architectures of the relay. Without frequency tracking, however, the f_S / f_{NOM} ratio is not necessarily an integer and the window length N is not a number of samples per cycle. Of course, we want N to be an integer. Moreover, to reduce the real-time processing burden – as we will explain later – we want N to be an even number. The value of N per (4) is an even number corresponding to a window length that is close to a nominal power system cycle.

In general, N does not have to be a multiple of a power system cycle. However, it is beneficial if N is close to a multiple of a power system cycle so that the h_D and h_Q filters notch out harmonics, at least for signal frequency that is close to nominal.

B. Orthogonal Filtering

The filter logic derives the orthogonal components of the input signal x using the h_D and h_Q orthogonal filters. We can optimize the computations by taking advantage of the symmetry of the h_D window and anti-symmetry of the h_Q window, as follows:

$$x_{D(n)} = \sum_{k=1}^{k=M/2} h_{D(k)}\cdot\left(x_{(n-M/2-k+1)} + x_{(n-M/2+k)}\right) \tag{5a}$$

$$x_{Q(n)} = \sum_{k=1}^{k=M/2} h_{Q(k)}\cdot\left(x_{(n-M/2-k+1)} - x_{(n-M/2+k)}\right) \tag{5b}$$

The numerical optimization in (5) is optional. It requires the present window length M to be an even number, i.e., the window grows by two samples with every two new samples of the input signal x.

The gains of the filters in (5) for the fundamental frequency f highly depend on the present window length M. Also, we want these gains to be exactly 1 only for the present system frequency f. To accomplish this, we correct the filters in (5) by using the multipliers C_D and C_Q presented in (8).

C. Gain Correction

Deriving the gain-correcting multipliers for the filters in (5) with the coefficients in (3) in the discrete time domain is complicated. Our implementation uses a high sampling rate of several kilohertz and allows us to greatly simplify derivation of the gain-correcting coefficients by assuming an infinite sampling rate (we derive the correcting coefficients in the continuous time domain).

In the continuous time domain, the gain of an FIR filter at frequency f can be calculated as an integral over a time equal to the window length of the product of the filter window function and the sine (or cosine) function of the frequency of interest f. Moreover, the gain does not depend on the phase alignment of the window function and the sine (or cosine) function. Therefore, we can select an arbitrary phase alignment that gives us the simplest integral to solve. We can also select either a sine or cosine function, depending on which function is easier to solve.

Following the above approach, we can obtain the continuous time-domain approximation of the gain coefficients as follows:

$$(C_D)^{-1} = \int_{-\frac{M}{N}\pi}^{\frac{M}{N}\pi} h_D(z)\cdot\cos\left(\frac{f}{f_{NOM}}\cdot z\right)dz \tag{6a}$$

$$(C_Q)^{-1} = \int_{-\frac{M}{N}\pi}^{\frac{M}{N}\pi} h_Q(z)\cdot\sin\left(\frac{f}{f_{NOM}}\cdot z\right)dz \tag{6b}$$

Equations (6) apply to any pair of orthogonal filters. For the orthogonal filters in (5), we write the following:

$$\left(C_D\right)^{-1} = \int_{-\frac{M}{N}\pi}^{\frac{M}{N}\pi} \cos(z)\cdot\cos\left(\frac{f}{f_{NOM}}\cdot z\right)dz \tag{7a}$$

$$\left(C_Q\right)^{-1} = \int_{-\frac{M}{N}\pi}^{\frac{M}{N}\pi} \sin(z)\cdot\sin\left(\frac{f}{f_{NOM}}\cdot z\right)dz \tag{7b}$$

Equations (7) are straightforward to solve and yield the following gain-correcting coefficients:

$$C_D = \left(\frac{M}{2}\cdot\left(\frac{\sin(A)}{A}+\frac{\sin(B)}{B}\right)\right)^{-1} \tag{8a}$$

$$C_Q = \left(\frac{M}{2}\cdot\left(\frac{\sin(A)}{A}-\frac{\sin(B)}{B}\right)\right)^{-1} \tag{8b}$$

where:

$$A = \pi\cdot\frac{M}{N}\cdot\left(\frac{f}{f_{NOM}}-1\right) \tag{8c}$$

$$B = \pi\cdot\frac{M}{N}\cdot\left(\frac{f}{f_{NOM}}+1\right) \tag{8d}$$

The value of A approaches 0 when the system operates near the nominal frequency. Of course, $\sin(A) / A$ in (8a) and (8b) approaches 1 if A approaches 0.

The equations in (8) show us that the gain-correcting coefficients depend on the per-unit system frequency f / f_{NOM}, the present per-unit window length M / N, and the present window length in samples M. System frequency does not change fast, and the f / f_{NOM} value can be refreshed relatively slowly. The rest of the operations involved in (8) can be implemented through a combination of real-time calculations and look-up tables.

As expected, the gain-correcting coefficients do not depend on the relay sampling frequency f_S because we derived these coefficients as approximations in the continuous time domain.

To evaluate accuracy of this simplified approach, we can calculate the true gain coefficients numerically in the discrete time domain for any given sampling frequency and compare these accurate values with the analytical approximations (8). Fig. 6 shows the approximated coefficients in (8) and the true coefficients for a 60 Hz nominal system frequency and sampling rates ranging from 2 to 10 kHz. The observed errors are very small: below a few percent for a sampling rate as low as 2 kHz and progressively smaller for higher sampling rates.

The values of the gain-correcting multipliers also inform us of how the filter works. When the window is short, the C_Q multiplier is large compared to the C_D multiplier, and the h_Q filter plays a bigger role. Because the quadrature filter (sine) is effectively a signal differentiator, it provides speed. When the window is longer, the C_Q multiplier decays to very small values and the h_D filter plays a bigger role. Because the direct filter (cosine) effectively averages the signal, it brings accuracy.

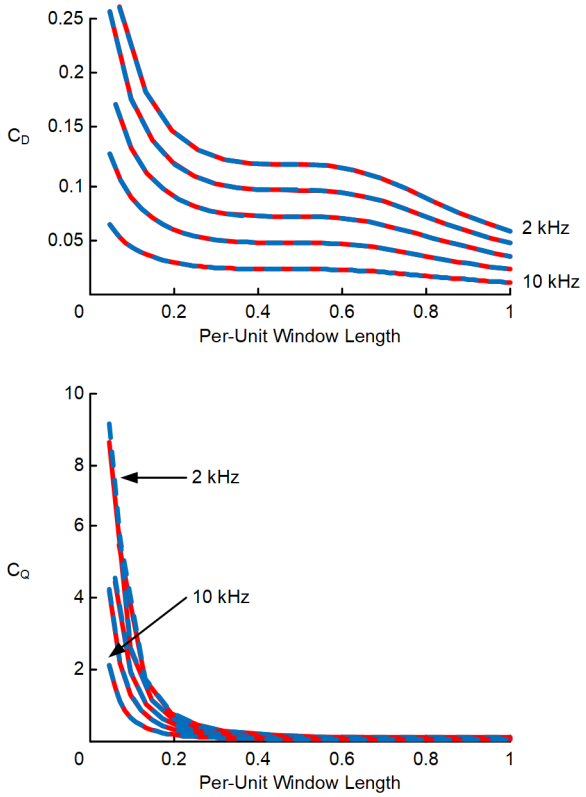

Fig. 6. Approximate (blue) and true (red) gain-correcting coefficients.

D. Group Delay Correction

The filter in Fig. 5 has a group delay of half the data window length M. Because M changes when the window is resized, the group delay is variable and it must be compensated for. When expressed in the frequency domain, this group delay calls for the following phase-shifting multiplier C_{DQ}:

$$C_{DQ} = 1\angle\left(\frac{\pi\cdot(M-0.5)}{N}\cdot\frac{f}{f_{NOM}}\right) \tag{9}$$

VII. Filter Operation Illustration With a Field Event

A relay [5] operated for an internal BG fault on a 345 kV, 109 mi line in a 60 Hz network; the network had a high level of series compensation located in the vicinity of the protected line. The relay recorded the voltages and currents shown in Fig. 7 and operated using a traveling-wave differential scheme, TW87 [3], in less than 2 ms. The relay actuated a two-cycle circuit breaker directly (without an interposing relay) by using a solid-state trip-rated output (10 µs closing time), and the breaker interrupted in 1.5 cycles. The fault lasted only 25 ms (1.75 cycles).

Fig. 8 shows the faulted phase current (raw and filtered). The steps in the plots mark the relay processing times. The relay samples its inputs at 1 MHz (not shown) for the application of traveling waves. It decimates the samples to 10 kHz for applications based on incremental quantities and further decimates the data to 2 kHz for applications based on phasors.

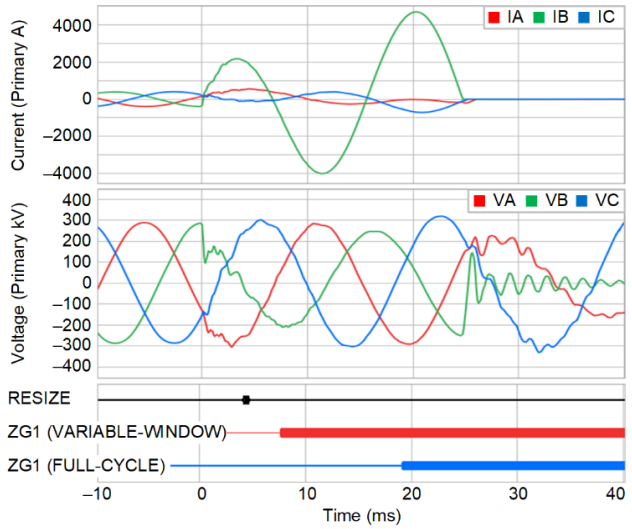

Fig. 7. Current and voltage signals during a field event.

The relay executes the described variable-window filter at the rate of 2 kHz. The figure shows a time lag between the raw inputs and the filtered outputs. This time lag is caused by additional low-pass filtering with an IIR filter prior to the variable-window filtering described in this paper.

To better illustrate the variable-window filter performance, Fig. 9 shows the faulted phase current and its magnitudes obtained with the presented method and, for comparison, with a full-cycle Fourier filter.

The variable-window filter logic resizes the window at about 4.5 ms into the fault (Fig. 7). Until that time, the filter is a full-cycle filter and it responds slowly to the fault current. When the window is resized, however, the filter output (Fig. 9 red trace) immediately settles around the true value of the input. With time, the window grows, and at about 20 ms into the fault, the window length is again one full cycle. From that time on, there is no difference between the variable-window filter and the reference full-cycle filter (Fig. 9 blue trace). From Fig. 8 and

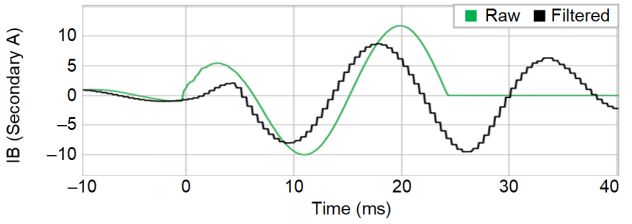

Fig. 8. Faulted phase current (raw and filtered).

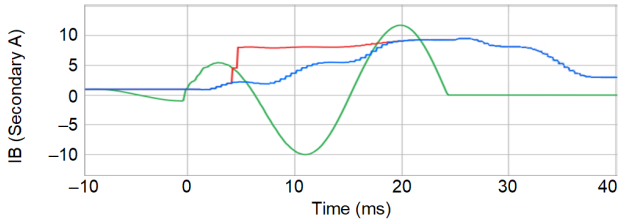

Fig. 9. Faulted phase current: raw (green) and magnitude measured with the presented method (red) and the full-cycle Fourier method (blue) for comparison.

Fig. 9, we see that the new algorithm does not show any overshoot and it settles within 1 ms after resizing.

Fig. 7 also shows the operation of the Zone 1 ground distance element for the fault. The reference full-cycle implementation responds in 19 ms, while the new distance algorithm, based on the variable-window filtering, responds in 7 ms, an operating time advantage of 12 ms or 0.7 cycle in the 60 Hz system.

VIII. Analysis and Discussion

A. Decaying DC Component

Our method uses sine-shaped and cosine-shaped base filters, so it is susceptible to errors (overshoot) due to the decaying dc component in the fault current. A mimic filter designed for the line X/R ratio is a well-known method to mitigate this problem. When used with a properly designed mimic filter, our method keeps the overshoot for current signals well below five percent, which is adequate for protection applications.

B. Low-Pass Prefiltering

Protective relays apply an analog anti-aliasing filter when digitizing the current and voltage inputs. Additionally, the relay designer may apply low-pass prefiltering as a method to improve security. These analog filters introduce a group delay and extend the transition from the predisturbance state to the disturbance state in the input to the variable-window filter. The resizing logic must wait an additional time before resizing the window if these filters are present. This delay ensures these analog filter artifacts do not pollute the short data window. An attempt to prefilter the input signal for better security adds a degree of delay (group delay). This delay must be added to the resizing logic to avoid filter artifacts. As a result, implementations with prefiltering effectively delay the output twice. Therefore, the overall design must be carefully optimized if intentional low-pass prefiltering is applied.

C. Off-Nominal Frequency Operation

The method fully and accurately compensates for off-nominal frequency operation, even though it uses a fixed sampling rate. Moreover, the method is well-suited for a wide range of off-nominal frequencies, such as for islanded power system operation with low-inertia machines. Our implementation uses the same sampling and processing rates for 50 and 60 Hz power systems and is accurate for frequencies between 40 and 70 Hz.

D. Harmonics

By using a filter window of a fixed length, the method notches out harmonics of the base frequency equal to f_S / N. This base frequency is very close to the nominal system frequency. Therefore, the filter effectively rejects harmonics of the nominal frequency. However, when the system frequency shifts away from the nominal value, the harmonic rejection is less effective. By comparison, FIR filters that use a variable sampling rate (frequency tracking) notch out harmonics completely, assuming they track the correct frequency. Therefore, the presented variable-window filter performs

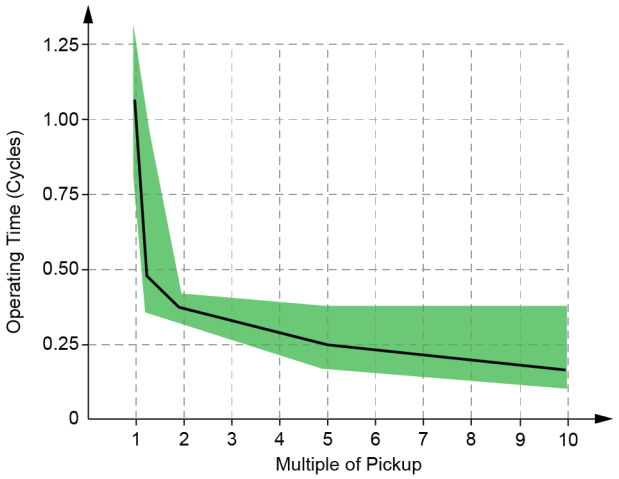

Fig. 10. Instantaneous overcurrent element operating time (range and median).

slightly worse with respect to harmonics than a frequency-tracking full-cycle filter. Nonetheless, it provides a degree of harmonic attenuation that is sufficient for protection applications.

E. Current Transformer Saturation

In most protection applications, current transformers (CTs) are sized to avoid saturation for at least the first full cycle after the fault. By using a short data window, the new filter allows fast operation before CTs saturate, and therefore the new filter improves dependability with respect to CT saturation. To illustrate this point, Fig. 10 shows the operating time of an instantaneous overcurrent element using the new filter. For multiples of pickup above 2, the element operates in less than a half cycle, including relay processing time. Therefore, these elements outrun CT saturation and operate dependably, even if the CT saturates after a half cycle. When CT saturation occurs, our method already uses a relatively long data window and it performs like any other full-cycle phasor estimator on the distorted secondary current waveforms.

F. Security

The new filter resizes the window in about a quarter cycle, and at that time it starts providing relatively accurate fault information to the downstream protection logic. At the time of resizing, however, the window length may be as short as one-eighth of a cycle and, as a result, the filter does not fully eliminate transient components. Therefore, we recommend that the downstream protection logic includes another quarter cycle for extra security when using the output from the variable-window filter. For example, the distance element described in more detail in the next section uses quarter-cycle coincidence timing for shaping the distance characteristic. If a protection element uses the quarter-cycle coincidence timing, when the element operates in about a half cycle, it uses a data window that is almost a half-cycle long but that entirely excludes the pre-fault-to-fault transition from the data window. Also, when it operates, the element has already been checking the operating conditions for about a quarter cycle by using relatively accurate inputs. This combination of removing the pre-fault-to-fault transition from the data window, using variable-window filtering, and applying additional quarter-cycle security in the downstream logic yields protection elements that are both consistently fast and secure.

IX. New Distance Element Design

Our algorithm applies to any comparator comprising a distance element (mho, reactance, blinders, directional, etc.). We describe the algorithm in relation to a general reach-sensitive comparator, such as the mho comparator or the reactance comparator. Fig. 11 shows the overall block diagram of the comparator logic, and the following subsections explain the key elements comprising the comparator logic.

A. Using Direct and Quadrature Components for Speed

For speed, our design applies coincidence timing to both the direct (real) and quadrature (imaginary) parts of the operating and polarizing signals. Depending on the point on wave (the moment of the fault as it relates to the peaks and zero-crossings of the pre-fault voltage), either the real part of a phasor or the imaginary part of a phasor develops faster. Typically, when the real part is slow, the imaginary part is faster, and vice versa. This speed relationship is caused by the fact that the real part is related to the signal value and the imaginary part is related to the signal derivative.

An identical coincidence timing logic, depicted in Fig. 2b, is applied separately to the real and imaginary parts of the operating and polarizing signals, with the outputs combined using OR gates, as Fig. 11 shows. Our solution applies the coincidence timer with the integrate-down option (see Fig. 3c).

B. Ensuring the Accuracy of Digital Coincidence Timing

Microprocessor-based relays, such as [5], apply high sampling rates and have enough processing power to implement time-domain comparators. However, unless the sampling and processing rates are very high, the time-domain comparator has limited steady-state accuracy. Assume a

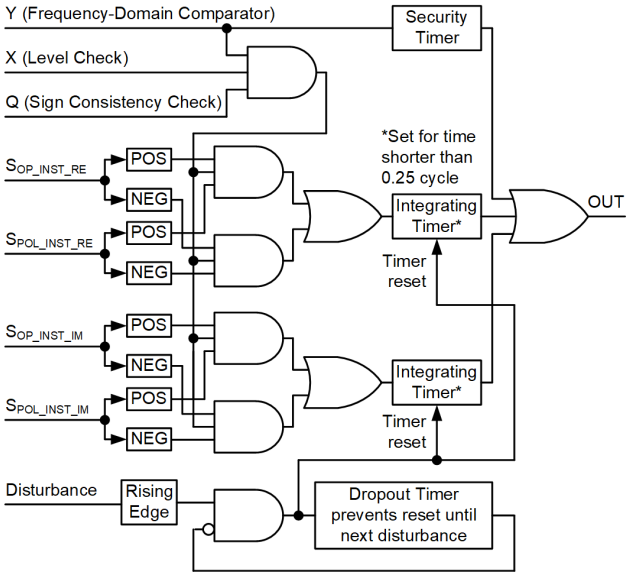

Fig. 11. New distance comparator logic diagram.

sampling rate of 2 kHz (a sampling period of 0.5 ms). In a 60 Hz system with a temporal resolution of 0.5 ms, the coincidence timer would have a comparator angle resolution of 360•0.5/16.67 = 10.8 degrees. This means that instead of the desired 90-degree comparator angle (4.17 ms coincidence timing), the logic performs either an 86.4-degree comparison (4 ms) or a 97.2-degree comparison (4.5 ms). The error can be reduced by higher sampling rates or by detecting changes in the polarity of the signals between the samples to accomplish subsample timing. Both these methods require more calculations and increase complexity.

Our design solves the accuracy problem by using a coincidence timer that is shorter than the accurate value. This yields a comparison angle greater than 90 degrees. Our design supervises such a time-domain comparator with a frequency-domain comparator (see Fig. 11). Because of the supervision, the final shape of the operating characteristic is equivalent to having an exact 90-degree comparator limit angle in the time-domain comparator. The frequency-domain comparator uses the same voltage and current phasor inputs and simply applies the logic of Fig. 1b.

C. Sign Consistency Check Between the Filtered and Raw Operating Signals

The raw and filtered (real part) operating signals are time-coherent because the filter compensates for the group delay. We can compare them sample by sample. In our design, we check the raw and filtered operating signals for consistent signs, i.e., if both are positive or negative. A disagreement in the signs tells us that the signals may have transients beyond the filtering capabilities of the applied filters. This is especially true when the variable-window filter uses very short data windows just after resizing or a CCVT creates transients that are large compared with the true operating signal. Fig. 12 illustrates the sign consistency logic. We use the output signal (Q) in Fig. 12 to supervise the AND gates in Fig. 11. When signal (Q) deasserts because the raw and filtered operating signals have opposite signs, the timers integrate down.

D. Checking the Level of the Raw Operating Signal

If the raw operating signal is small, an internal or external fault is very close to the reach point. To add margin for transients for such faults, the element applies more security when operating for very small levels of the raw operating signal. Fig. 13 illustrates the operating signal level logic. We use the output signal (X) in Fig. 13 to supervise the AND gates in Fig. 11. A low signal level, such as below one percent of the nominal voltage, causes the timers to integrate down.

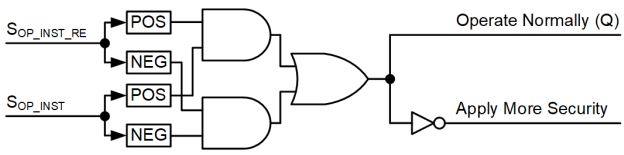

Fig. 12. Sign consistency logic checking the raw (S_{OP_INST}) and filtered ($S_{OP\_INST\_RE}$) operating signals.

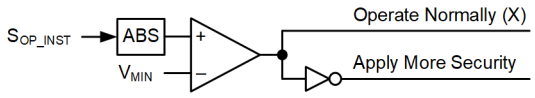

Fig. 13. Raw operating signal level logic.

E. Resetting Coincidence Timers Upon Disturbance

There may be conditions when the comparator is asserted prior to a fault. This is typically the case for the reactance comparator. The load impedance is typically below the reactance set point. If so, the load asserts the output of the reactance comparator. A mho comparator under heavy load conditions may assert as well (the mho element is not operating on load because it is typically blocked by the load-encroachment logic, but the mho comparator itself may be permanently asserted on load).

If the mho or reactance comparator is permanently asserted on load, it has a lower security margin for a subsequent fault external to the zone of distance protection. With reference to Fig. 11, our design uses a disturbance detector to reset the integrating timers. This way, the timers forget their memory of the pre-fault load and start fresh, using only the fault data.

F. Dependability for Very Small Operating Signals

By design (see Section IX.D), the fast coincidence timing algorithm restrains if the operating signal is too small. To maintain steady-state accuracy and dependability, our design uses the frequency-domain comparator (Fig. 1b) with a time delay on the order of one to two cycles (see Fig. 11).

G. Zone 1 Dynamic Reach

The Zone 1 distance element is normally set to underreach the remote line terminal and trip directly without the pilot channel. To improve security, our design dynamically reduces the Zone 1 reach to about 80 percent of the set reach when the filter resizes. Subsequently, the Zone 1 reach grows with the filter window length, and it reaches 100 percent of the set value when the filter window reaches one full cycle.

X. Laboratory Test Results

The presented distance element design has been implemented on a relay platform based on [5] and tested for security, dependability, and operating times under a variety of system conditions using a real-time digital simulator (RTDS). The Zone 1 distance elements are applied to underreach the remote line terminal to trip directly without a pilot channel. The Zone 2 distance elements are applied to overreach the remote line terminal as part of a pilot scheme or for step distance protection. Therefore, our design applies the solution described in Section IX to Zone 1, and for Zone 2, it uses a simplified design biased for speed without strict transient reach accuracy requirements. As a result, the operating times and transient accuracy of Zones 1 and 2 differ.

Fig. 14 shows the Zone 1 operating time curves as a function of fault location, respective to the set reach, for a range of source-to-impedance ratios (SIRs). In strong systems (SIR of 0.1), the element operates in less than half a cycle for locations up to about 80 percent of the set reach. In weaker systems (SIR

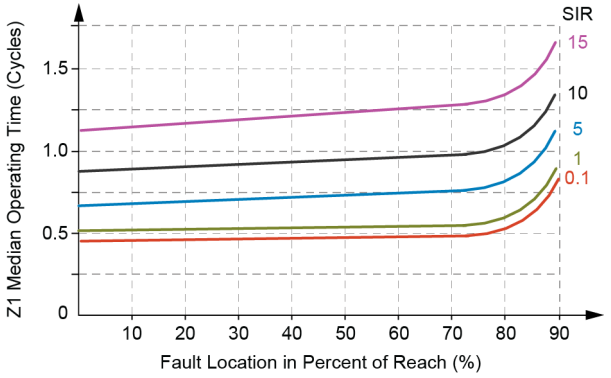

Fig. 14. Underreaching Zone 1 operating times.

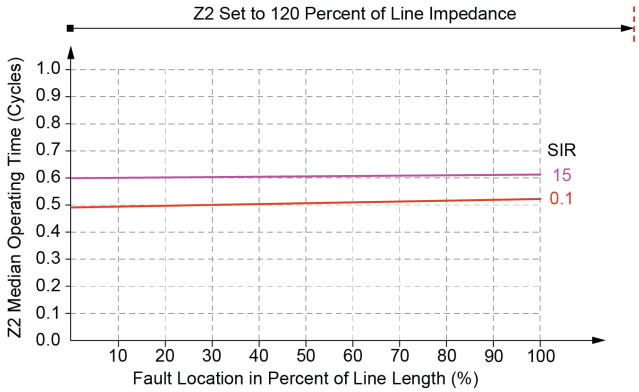

Fig. 15. Overreaching Zone 2 operating times.

of 15), the element operates as fast as 1.2 cycles. The Zone 1 element has excellent transient accuracy with a transient overreach below five percent. The operating time curves bend up for locations 20 percent short of the reach point as a result of the solutions described in Sections IX.D and IX.G.

Fig. 15 shows the Zone 2 operating time curves. The zone is set to 120 percent of the line, and it provides a near-constant operating time for faults anywhere along the line. Having relaxed transient overshoot requirements, Zone 2 is slightly faster than Zone 1.

Fig. 14 and Fig. 15 apply to both the mho and quadrilateral distance elements (in our design, there are very small differences in speed and transient overreach between the two types of distance elements).

XI. Distance Element Illustration With a Field Event

Consider the field case included in Section VII (reprinted in Fig. 16 with operating signals). In our tests using this field recording, the variable-window filter resized the window at about 4.5 ms into the fault. The Zone 2 mho and quadrilateral elements, set to 120 percent of the line impedance, responded in about 7.6 ms. The Zone 1 mho and quadrilateral elements, set to 85 percent of the line impedance, responded in about 8 ms. These operating times include the relay processing time and the trip-rated solid-state output closure time.

Fig. 17 illustrates operation of the Zone 1 mho elements by showing the real and imaginary parts of the polarizing and operating signals. Before the fault, the polarizing and operating signals are out-of-phase. When the fault occurs, the operating

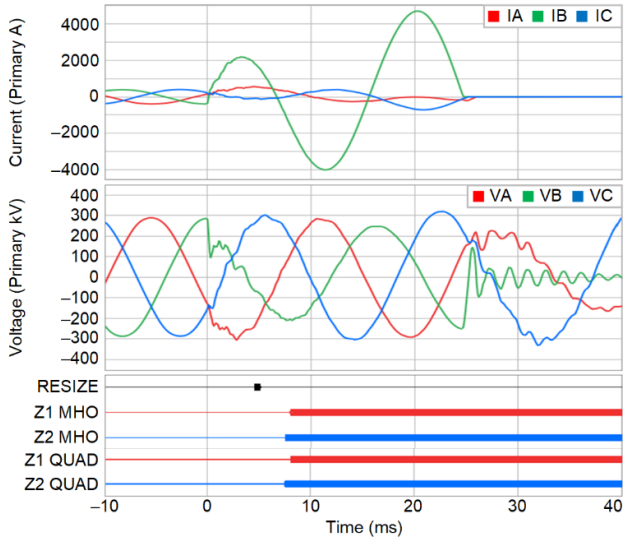

Fig. 16. Current and voltage signals during a field event.

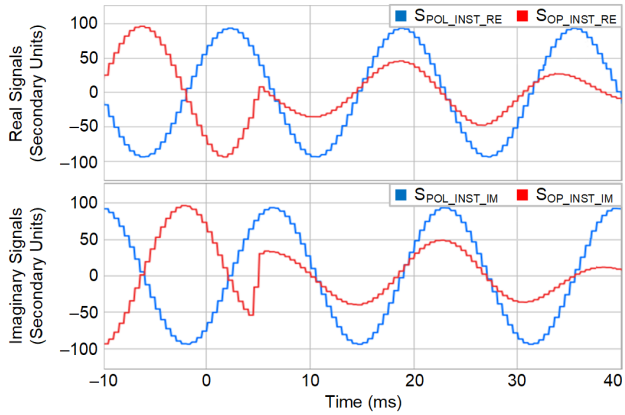

Fig. 17. Zone 1 mho operating and polarizing signals (real and imaginary parts).

signal begins changing, and when the filter window resizes, it jumps to reflect the fault value. About 4.5 ms into the fault, the imaginary parts of the polarizing and operating signals are the same polarity, which engages the coincidence timer and results in Zone 1 operation. The real parts are the same polarity starting at about 7 ms. S_{OP_INST} (I•Z – V signal) is large (see Section IX.D) and very clean, despite distortions in the relay voltages and currents. The Zone 1 operation in 8 ms is robust and secure. The polarizing signal in Fig. 17 does not change during the fault because the element is fully memory-polarized.

Fig. 18 illustrates the part of the comparator logic that determines if the raw and filtered operating signals have the same polarity. The two signals agree very well (except during the time interval between 0 and 4.5 ms, i.e., before the window

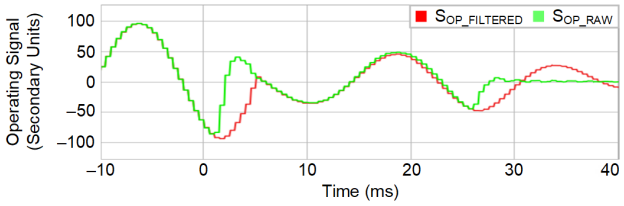

Fig. 18. Zone 1 mho raw and filtered operating signals.

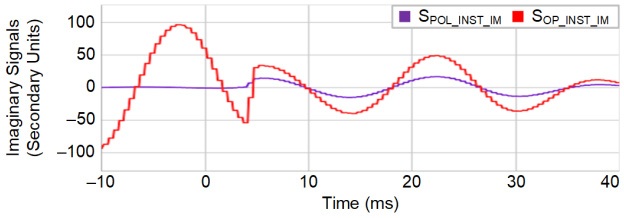

Fig. 19. Zone 1 quadrilateral operating and polarizing signals (imaginary parts).

resizing), which allows the distance element to operate fast (see Section IX.C).

Fig. 19 illustrates the operation of the Zone 1 quadrilateral elements by showing the imaginary parts of the polarizing and operating signals (this example shows quadrilateral elements that are polarized with the loop current). Prior to the fault, the polarizing and operating signals are out-of-phase. When the fault occurs, the operating and polarizing signals begin to change, and when the filter window resizes at about 4.5 ms, both jump to reflect the fault value. About 5 ms into the fault, the imaginary parts of the polarizing and operating signals are the same polarity, which engages the coincidence timer and results in Zone 1 operation.

XII. Conclusion

This paper presents the implementation of a distance protection element in a microprocessor-based relay that uses coincidence timing and window resizing. The paper explains the principles of coincidence timing and its benefits. Taking advantage of digital technology, the paper introduces several enhancements complementing the classic coincidence timing method. The method presented in this paper achieves an excellent balance between speed (half-cycle operating time in strong systems and one-cycle operating time in weak systems) and security (Zone 1 transient overreach below five percent).

The new distance element design uses the concept of variable-window filtering. The filter uses an explicit resizing logic with several security conditions to allow window resizing only when it is secure to do so. The logic intentionally delays window resizing to ensure that the short data window only includes disturbance samples and excludes the pre-fault-to-fault transition data. The paper derives the filter for the sine- and cosine-shaped base filters, and it teaches how to design the filter for any pair of orthogonal filters using the continuous time-domain approximation method. The described filter is fully compensated for off-nominal frequencies, and it rejects harmonics reasonably well even if the frequency deviates from the nominal value.

The distance element design described in this paper is intended for relay hardware with fixed sampling and processing rates. We prefer these relay architectures for their internal simplicity, especially when implementing time-domain protection principles. The distance element design presented in this paper has been implemented in a relay platform based on [5], and it operates consistently with trip times of a half cycle, including relay processing time and trip-rated output contact closure time.

XIII. References

[1] B. Kasztenny, M. V. Mynam, T. Joshi, A. Kathe, and C. Daniels, "A New Digital Distance Element Implementation Using Coincidence Timing," proceedings of the 15th International Conference on Developments in Power System Protection, Liverpool, UK, March 2020.

[2] B. Kasztenny, M. V. Mynam, T. Joshi, and C. Daniels, "A New Digital Filter Using Window Resizing for Protective Relay Applications," proceedings of the 15th International Conference on Developments in Power System Protection, Liverpool, UK, March 2020.

[3] E. O. Schweitzer, III, B. Kasztenny, A. Guzmán, V. Skendzic, and M. V. Mynam, "Speed of Line Protection – Can We Break Free of Phasor Limitations?" proceedings of the 68th Annual Conference for Protective Relay Engineers, College Station, TX, March 2015.

[4] B. Kasztenny, "A New Method for Fast Frequency Measurement for Protection Applications," proceedings of the 13th International Conference on Developments in Power System Protection, Edinburgh, UK, March 2016.

[5] SEL-T400L Instruction Manual. Available: https://selinc.com.

XIV. Biographies

Bogdan Kasztenny has over 30 years of experience in power system protection and control. In his decade-long academic career (1989–99), Dr. Kasztenny taught power system and digital signal processing courses at several universities and conducted applied research for several relay manufacturers. In 1999, Bogdan left academia for relay manufacturers where he has since designed, applied, and supported protection, control, and fault-locating products with their global installed base counted in thousands of installations. Bogdan is an IEEE Fellow, a Senior Fulbright Fellow, a Distinguished CIGRE Member, and a registered professional engineer in the province of Ontario. Bogdan has served as a Canadian representative of the CIGRE Study Committee B5 (2013–2020) and on the Western Protective Relay Conference Program Committee (2011–2020). In 2019, Bogdan received the IEEE Canada P. D. Ziogas Electric Power Award. Bogdan earned both the Ph.D. (1992) and D.Sc. (Dr. habil., 2019) degrees, has authored over 220 technical papers, and holds over 50 U.S. patents.

Mangapathirao (Venkat) Mynam received his MSEE from the University of Idaho in 2003 and his BE in electrical and electronics engineering from Andhra University College of Engineering in 2000. He joined Schweitzer Engineering Laboratories, Inc. (SEL) in 2003 as an associate protection engineer in the Engineering Services division. He is presently working as an engineering director in SEL Research and Development. He was selected to participate in the U. S. National Academy of Engineering (NAE) 15th Annual U. S. Frontiers of Engineering Symposium. He is a senior member of IEEE and holds patents in the areas of power system protection, control, and fault location.

Chad Daniels received a bachelor's degree in computer science from Eastern Washington University in 2006. He spent a decade working with data compression and error-correction technologies before joining Schweitzer Engineering Laboratories, Inc. (SEL) in 2016. He is presently working in SEL Research and Development where he writes firmware for ultra-high-speed microprocessor-based relays.

Titiksha Joshi is a lead power engineer in Research and Development at Schweitzer Engineering Laboratories, Inc. She received her bachelor's degree in electrical engineering from the University of Mumbai in 2012 and a master of science degree in electrical engineering from Arizona State University in 2014. She has worked as an intern at Crompton Greaves (2011) and Midcontinent Independent System Operator, Inc. (2013).

Amol Kathe received a bachelor's degree in electrical engineering from the University of Mumbai in 2006 and a master's degree in electrical engineering from Michigan Technological University in 2014. He worked for over five years as an operations manager in the Energy department at Reliance Infrastructure Limited. He joined Schweitzer Engineering Laboratories, Inc. in 2014 and holds the position of lead power engineer in Research and Development.