Observations on the Application of the IEEE C37.114 Double Ended Fault Location Method

Introduction

At the 2013 Georgia Tech Fault and Disturbance Analysis Conference, Dominion Virginia Power and Softstuf presented on the application of the IEEE C37.114 Double Ended Fault Location Method after months of collaboration. That paper and presentation detailed the team's findings on applying the Double Ended (DE) method to real system faults on Dominion's electric grid, developing programs and software tools to apply the method, and practices on selecting the best data to use in the method.

In the two years since that paper was published, the Dominion and Softstuf team has continued to collaborate on the DE method. One of the main goals in this timeframe was to automate the application of the DE method, allowing computers and servers to calculate DE fault location results from any fault data available from the electric grid. This type of automation is desired to allow rapid analysis and decision making abilities, ultimately to help restore electric service to customers and transmission equipment as fast as possible.

This quest for automation has led to many breakthroughs in implementation of the DE fault location method, which will be shared in this paper. This paper will present the findings and results from actual faults on the Dominion system. The team has found through these real system cases where the DE method doesn't perform according to initial expectations. Comparisons and study of the real fault data have lead to determining new best practices and realizations, all of which will be shared in this paper.

Through these best practices and findings by the Dominion and Softstuf team on the DE method, significant improvements were made in the application of the DE method that led to the successful automation of the Double Ended fault location method.

The Mathematical Approach

The C37.114 double ended method uses sequence components to transpose an unbalanced three phase system to a vector sum of two balanced systems (positive and negative sequences) and an offset system (zero sequence). For example, an unbalanced three phase system with voltage vectors $\mathbf{V_A}$, $\mathbf{V_B}$, $\mathbf{V_C}$ is shown in Figure 1. The resulting zero, positive, and negative sequence systems are shown in Figure 2 (superimposed with the unbalanced system).

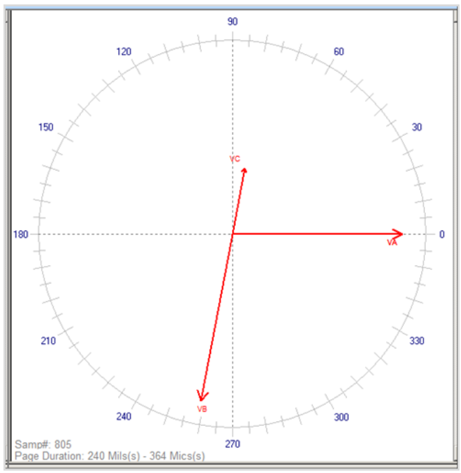

Figure 1: Three phase voltage vectors in an unbalanced state

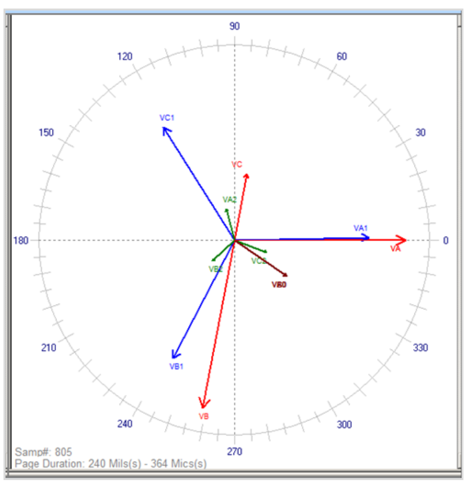

Figure 2: Resulting Zero, Positive and Negative Sequence Systems
(Brown, Blue and Green Respectively)

As seen in Figure 2, the positive sequence vectors $\mathbf{V_{A1}}$, $\mathbf{V_{B1}}$, $\mathbf{V_{C1}}$ are balanced with equal magnitudes and 120 degrees angles. The negative sequence vectors $\mathbf{V_{A2}}$, $\mathbf{V_{B2}}$, $\mathbf{V_{C2}}$ are also balanced. The zero sequence vectors $\mathbf{V_{A0}}$, $\mathbf{V_{B0}}$, $\mathbf{V_{C0}}$ are equal offsets (shown displayed over each other).

The mathematical relationships between the unbalanced system and the sequence components systems are shown in the following vector sum:

$\mathbf{V_A = V_{A0} + V_{A1} + V_{A2}}$

$V_B = V_{B0} + V_{B1} + V_{B2}$
$V_C = V_{C0} + V_{C1} + V_{C2}$

Similarly, for unbalanced three phase currents I_A, I_B, I_C we get the following vector sum relationships:
$I_A = I_{A0} + I_{A1} + I_{A2}$
$I_B = I_{B0} + I_{B1} + I_{B2}$
$I_C = I_{C0} + I_{C1} + I_{C2}$

The sequence components vectors are calculated as a shifted sum of the unbalanced vectors. The sequence vectors for phase-A (meaning with phase-A as the reference phase) are:
$V_{A0} = ⅓ (V_A + V_B + V_C)$
$V_{A1} = ⅓ (V_A + (V_B * 1 \lfloor\underline{120}) + (V_C * 1 \lfloor\underline{-120}))$
$V_{A2} = ⅓ (V_A + (V_B * 1 \lfloor\underline{-120}) + (V_C * 1 \lfloor\underline{120}))$
$I_{A0} = ⅓ (I_A + I_B + I_C)$
$I_{A1} = ⅓ (I_A + (I_B * 1 \lfloor\underline{120}) + (I_C * 1 \lfloor\underline{-120}))$
$I_{A2} = ⅓ (I_A + (I_B * 1 \lfloor\underline{-120}) + (I_C * 1 \lfloor\underline{120}))$

The sequence vectors for phase-B are:
$V_{B0} = ⅓ (V_A + V_B + V_C)$
$V_{B1} = ⅓ ((V_A * 1 \lfloor\underline{-120}) + V_B + (V_C * 1 \lfloor\underline{120}))$
$V_{B2} = ⅓ ((V_A * 1 \lfloor\underline{120}) + V_B + (V_C * 1 \lfloor\underline{-120}))$
$I_{B0} = ⅓ (I_A + I_B + I_C)$
$I_{B1} = ⅓ ((I_A * 1 \lfloor\underline{-120}) + I_B + (I_C * 1 \lfloor\underline{120}))$
$I_{B2} = ⅓ ((I_A * 1 \lfloor\underline{120}) + I_B + (I_C * 1 \lfloor\underline{-120}))$

And finally, the sequence vectors for phase-C are:
$V_{C0} = ⅓ (V_A + V_B + V_C)$
$V_{C1} = ⅓ ((V_A * 1 \lfloor\underline{120}) + (V_B * 1 \lfloor\underline{-120}) + V_C)$
$V_{C2} = ⅓ ((V_A * 1 \lfloor\underline{-120}) + (V_B * 1 \lfloor\underline{120}) + V_C)$
$I_{C0} = ⅓ (I_A + I_B + I_C)$
$I_{C1} = ⅓ ((I_A * 1 \lfloor\underline{120}) + (I_B * 1 \lfloor\underline{-120}) + I_C)$
$I_{C2} = ⅓ ((I_A * 1 \lfloor\underline{-120}) + (I_B * 1 \lfloor\underline{120}) + I_C)$

The concept of the sequence components method is to use the vectors from one of the sequence systems instead of using those of the unbalanced system. This provides a formidable reduction in processing requirements. For example, if we choose the phase-A negative sequence system then instead of having to work with V_A, V_B, V_C, and I_A, I_B, I_C we just work with V_{A2} and I_{A2}. Similarly if we choose the phase-B zero sequence system then we just work with V_{B0} and I_{B0}.

Accordingly, the accuracy of the sequence components method is based on two choices: 1) which reference phase to use, and 2) under which sequence components system. With three phases and three components per phase, we get a total of nine combinations. The choice of which combination to use depends on the fault characteristics as presented in the Real World Data Approach section.

A number of observations about the mathematical structure are worth noting:

1) Choosing the reference phase is important. It is best to use an unfaulted phase as the reference.

2) Choosing the sequence components system is also important. The fault type must contain appreciable values of the used sequence components to provide reliable locations. For instance, while zero and negative sequence can work well for ground faults, balanced 3 phase faults, having minimal negative and zero sequence components, require positive sequence for dependable results.

3) Precise time synchronization is not required because the math is vector based. So long as a valid set of fault vectors are measured from each end of the line then the equations can be used regardless if the measurements are cycles apart.

Perform the calculation twice

One might think simply performing the double ended calculation would be sufficient to provide a reasonable location, but we have found that performing the calculation twice, reversing the near side and far side quantities, provides valuable insight into the accuracy of the result. A network diagram and the equation are shown below.

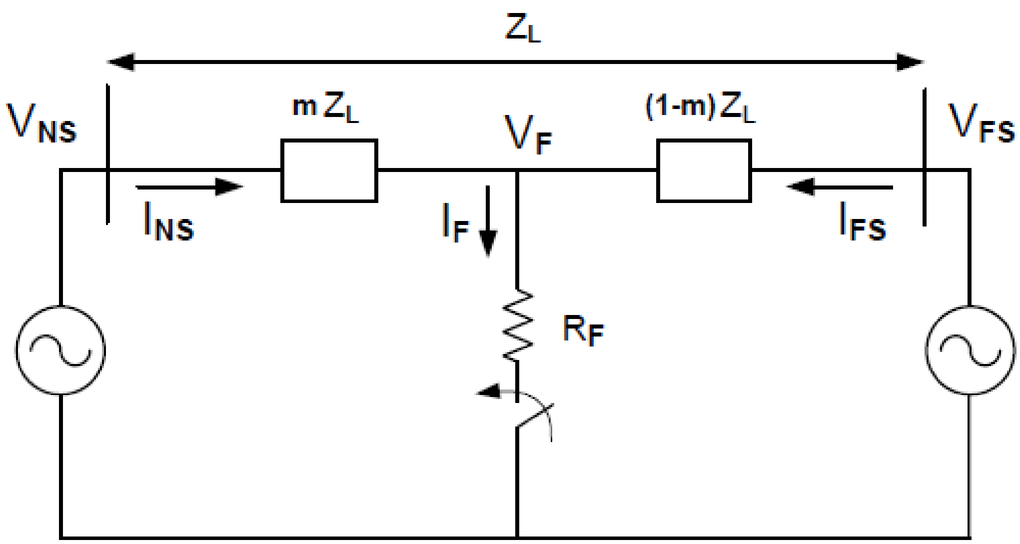

Figure 3: Double ended network diagram

From the diagram shown in Figure 3 we can derive the equation:

$$m = (V_{NS} - V_{FS} + Z_L I_{FS}) / Z_L (I_{NS} + I_{FS})$$

Under ideal conditions, this equation would yield identical results when swapping the quantities from the near side and far side. In practice, however, we have found this is not always the case. If the results exactly agree on a location, then we have a higher confidence in the results. If the results are notably different, then we will question the results. When results differ it almost always results in “overlap” in the target fault location. We call it overlap because the results often are too long and do not add up to the actual length of the transmission line. This may seem obvious since overlap inherently indicates that the locations do not agree, but this information would not be known if the calculation is only performed once. Visualizing the overlap, or lack thereof, provides a quick indication of the quality of the result. Figure 4 shows an example of a result with no overlap, while Figure 5 shows overlapping fault locations.

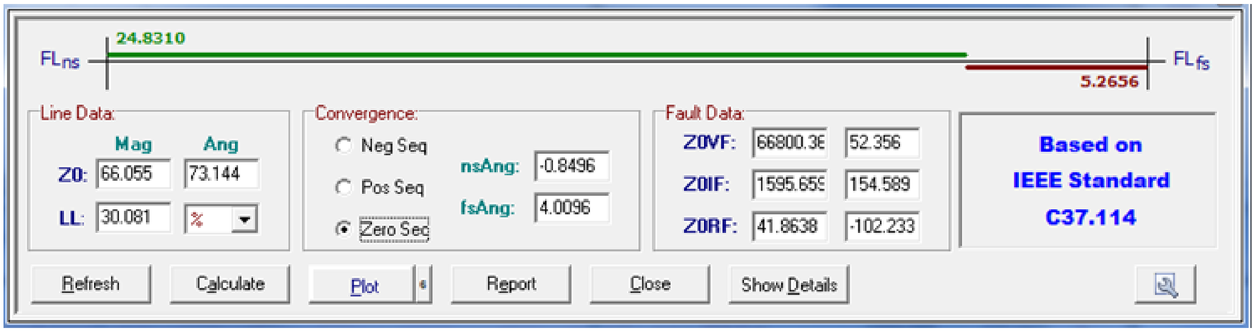

Figure 4 - No overlap in results

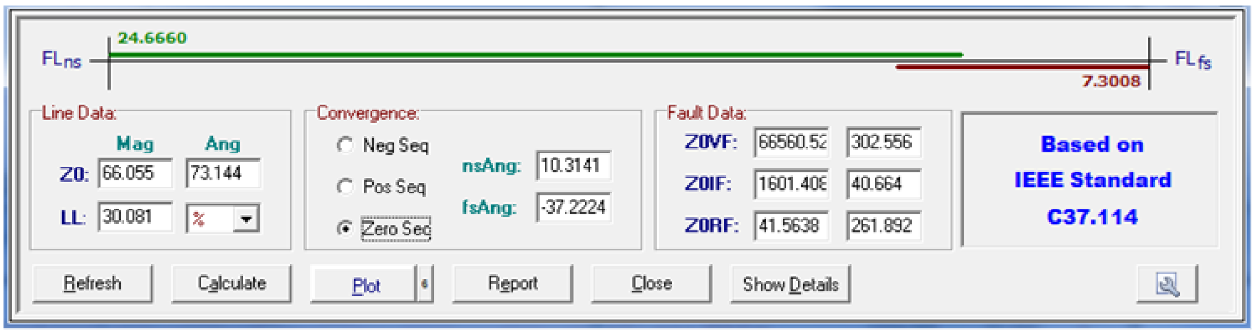

Figure 5 – Overlap in results

One other variable is the angle of the result. The double ended equation uses complex numbers as the input quantities and therefore the output is also a complex number, usually expressed as a magnitude and an angle. One would expect that a fault location on a transmission line would have no angle, or a zero angle, but that is not always the case. When performing the calculation twice, swapping near and far sides, we can end up with a variety of possible results, such as a low angle at one end and a large angle at the other. Through our experience we have found resultant angles of 15 degrees or less at **both ends** usually indicate an accurate answer.

Real World Data Approach

Fifteen faults, five per voltage level at 500kV, 230kV, and 115kV were evaluated to determine the most effective approach for double ended analysis. Events chosen were those we had a known cause and a field-verified fault location. Faults were predominantly phase to ground, but did include one two-phase to ground fault and one three-phase fault. Each fault event was tested in 18 ways, comparing positive, negative, and zero sequence components against each phase reference, using two cursor points within the fault. The measurements recorded are shown in Table 1. "Vref" in Table 1 refers to the reference voltage used when calculating the symmetrical components.

Cursor	Stable			1.75 Cycles		
Sequence	Positive	Negative	Zero	Positive	Negative	Zero
Vref	A	A	A	A	A	A
	B	B	B	B	B	B
	C	C	C	C	C	C

Table 1: Measurements from each fault

In Table 1, "Cursor" refers to the point in the fault at which the double ended calculation was made. "Stable" indicates the point was chosen based on a visual inspection of a single-ended trend of fault stability. "1.75 Cycles" indicates the point used was 1.75 cycles after fault inception time. Figures 6 and 7 show cursor placement for the two methods.

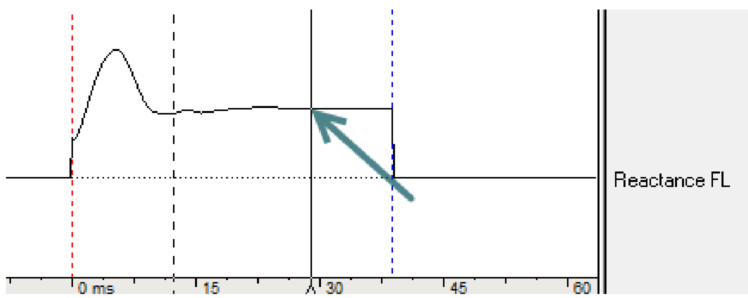

Figure 6: Stable Cursor Location

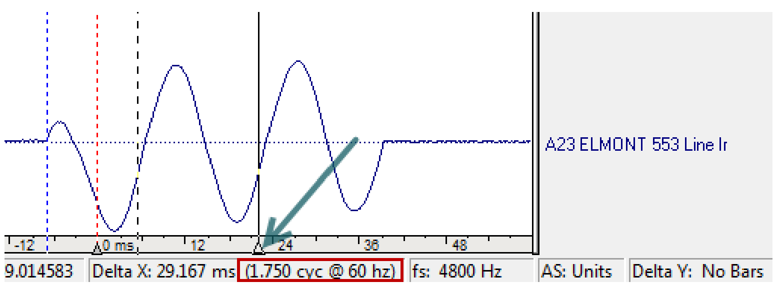

Figure 7: 1.75 Cycles Cursor Location

With each test, a near-side and far-side calculation was made, yielding two locations per test. Each terminal is then given an accuracy rating indicating their calculated location's relative position to the known location as a percentage beyond or short of the actual fault location. The terminal ratings are then combined into an overall accuracy assessment for the entire test by averaging the absolute values of each result's deviation percentage from the known location.

The tests also calculate an aforementioned "overlap" value indicating what portion of the line is encompassed by both locations. If one terminal's location is beyond the remote terminal's location, you get a positive overlap for the portion of the line between the two locations. A negative overlap implies a gap between the two locations from each side. The overlap value is defined as the percentage of the line existing within the overlap or within the gap.

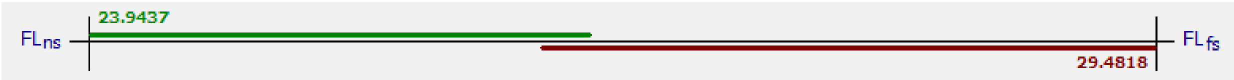

Figure 8: Example of Positive Overlap

The angle portion of the double ended result is being considered as well. An angle “accuracy” value is determined by averaging each side’s deviation from zero degrees. It averages the absolute values of each result’s angle, normalized between -180 and 180 degrees. This value is used to correlate angular deviation to overall accuracy.

Each of the fifteen faults also received an inferred resistivity rating based on the relative angles between current and voltage on the faulted phase(s). Tighter angles should associate with resistive faults, while angles more closely matching the line angle are assumed to be low impedance faults.

Results

Cursor Location

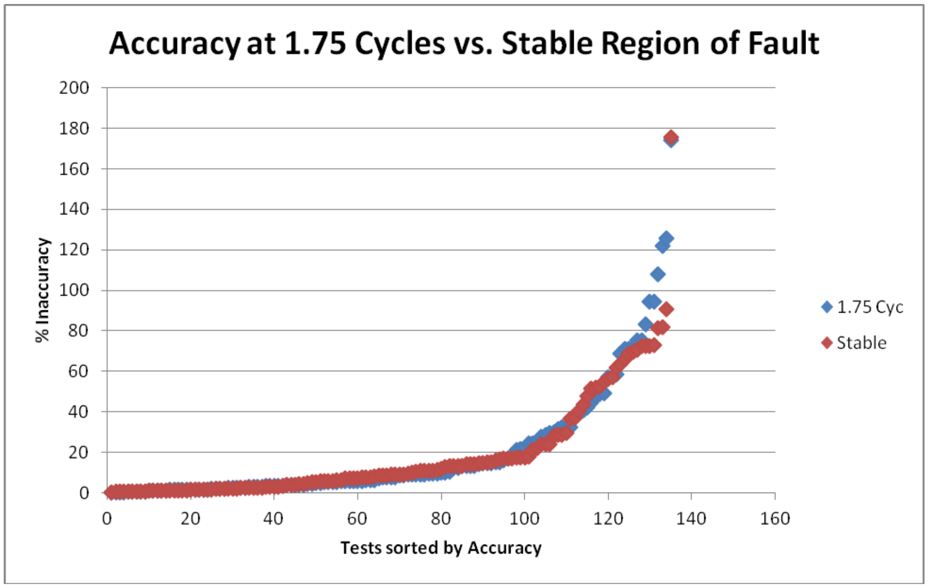

Graph 1

Graph 1 compares the accuracy of the two cursor location methods used in our tests. Results are sorted by test accuracy. Based on our data, there is little perceivable advantage to using a stable region of the fault. Both methods trend very similarly. An argument could be made for slightly higher accuracy when using stability toward the far right of the graph, but neither result in that region is yielding an accurate location regardless. This also corroborates our finding that precise time synchronization is unnecessary. The stability approach in particular makes no effort to synchronize cursor points between the terminals, and despite this, it is capable of providing very good results.

Voltage level and sequence components used have no notable impact on the cursor region results.

Sequence Components

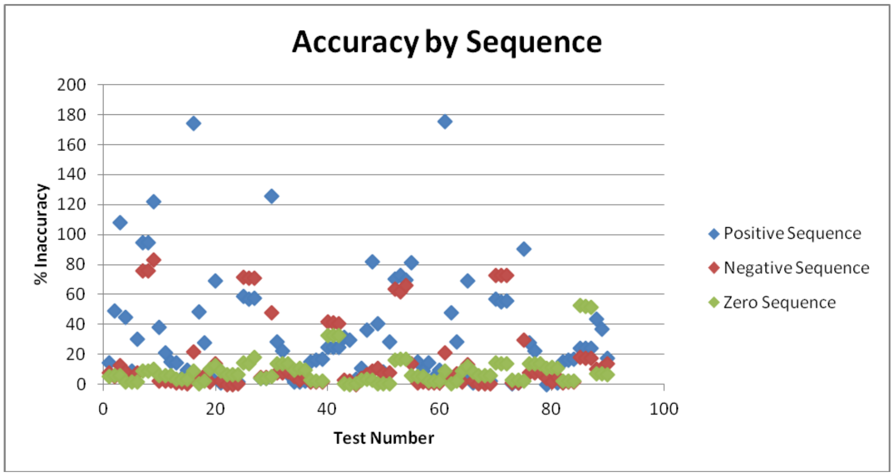

Graph 2

Graph 2 compares double ended result accuracy analysis across positive, negative, and zero sequence components. Positive sequence components generally yield the worst results. Negative and zero sequence average much better, with negative being the most likely to be very accurate. Zero sequence is the least prone to being wildly off, however, nearly always remaining within 20% accuracy.

Positive sequence was found to be especially inaccurate at 500kV. Otherwise no significant trends could be seen across voltage level.

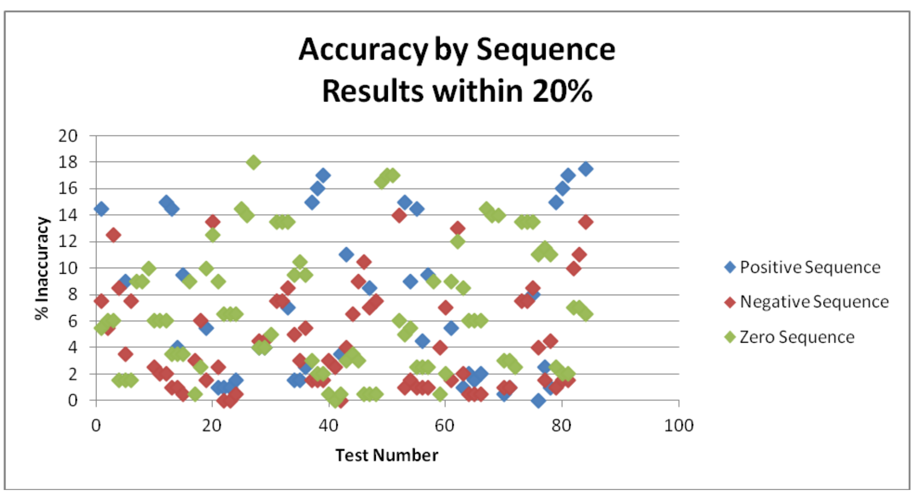

Graph 3

Graph 3 shows the same spread of data, although limited to results with 20% accuracy or better. While the trend is less defined, negative sequence can be seen as the most frequent within 5% or better, implying the best precision when correct. This can also be observed in tabular format, seen below.

% Accuracy	Overall			Within 20% Accuracy		
	Positive	**Negative**	**Zero**	**Positive**	**Negative**	**Zero**
Max	175.50	83.00	**52.50**	N/A	N/A	N/A
Average	33.10	**16.19**	**9.21**	7.23	**5.03**	**6.85**
Median	22.00	**5.75**	**6.00**	5.50	**3.50**	**6.00**

Table 2

Table 2 provides the max, average, and median values for our postive, negative, and zero sequence results, including both overall results and results filtered down to within 20% accuracy. Zero sequence shows the best overall average, markedly better than the negative sequence overall average. By contrast, however, the negative sequence overall median is slightly better. This juxtaposition implies a higher likelihood of precision with negative sequence despite a worse overall average accuracy. The negative sequence median advantage is more pronounced when only considering results with better than 20% accuracy, a range in which negative sequence also has the best average. Again, the implication is that if not for those negative sequence results that are wildly off, it would generally provide the best location. It should be noted zero sequence had the most tests coming in under 20% accuracy, with 84 results. Negative and positive had 71 and 43 results respectively.

Given the above findings, we may infer the following: If negative and zero sequence largely agree, negative sequence should be trusted as the most accurate. If there is significant disagreement, zero is likely to provide the better location, although it’s unlikely to be dead on.

Implications of Overlap and Double Ended Result Angle

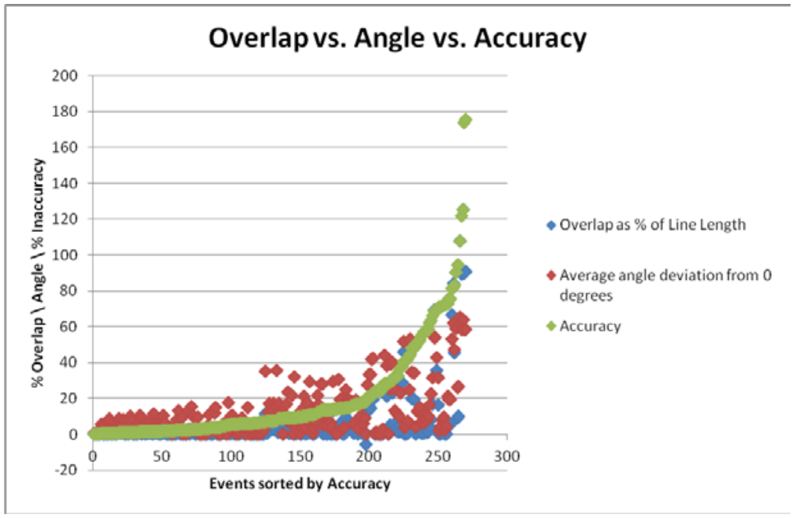

Graph 4

Graph 4 shows the relationship between accuracy, average double ended result angle deviation, and overlap. An expected correlation can be drawn between increased overlap and decreased accuracy. Significant overlap inherently results in poor accuracy, as the two terminals by definition disagree on location. Average DE fault angle deviation from zero degrees also shows a notable correlation to accuracy, with tighter angles clearly associating with better accuracy. A slight variation in angle seems acceptable, however, with many good results being generated at angles less than 15 degrees. Performance across voltage levels was fairly consistent.

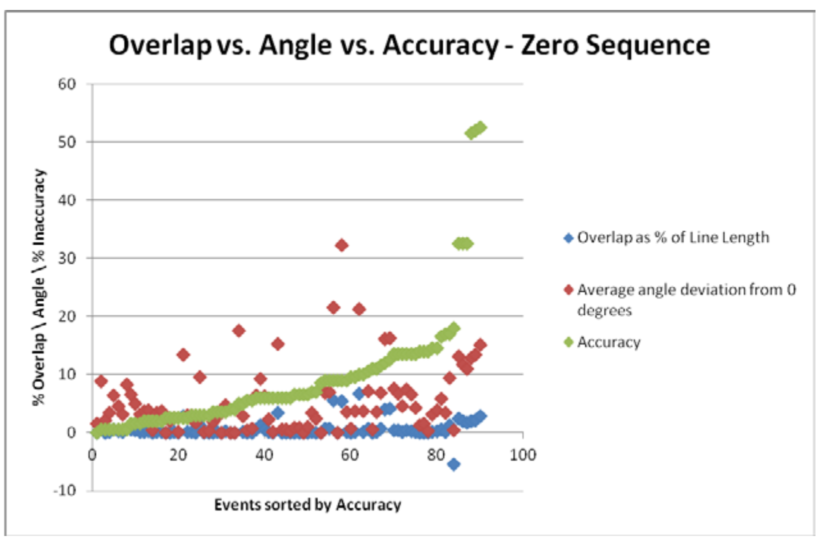

Graph 5

Interestingly, zero sequence quantities resulted in very little overlap, distinctly less than negative and positive sequence, as displayed in Graph 5. As discussed previously, zero sequence may not be the most accurate, but it does typically show both terminals as agreeing on location. This may even be considered detrimental, as it gives a false sense of certainty compared to negative sequence evaluation.

Impact of Referencing Faulted Phases

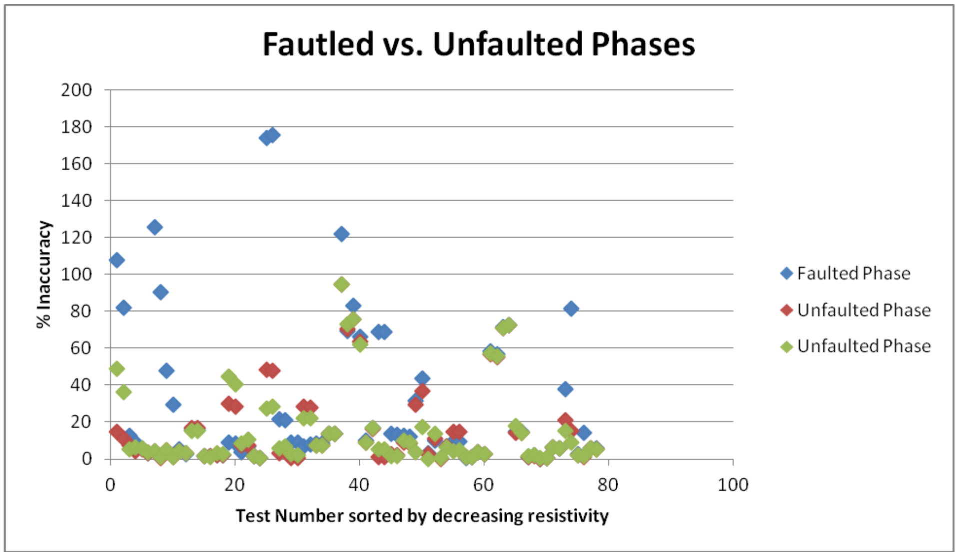

Graph 6

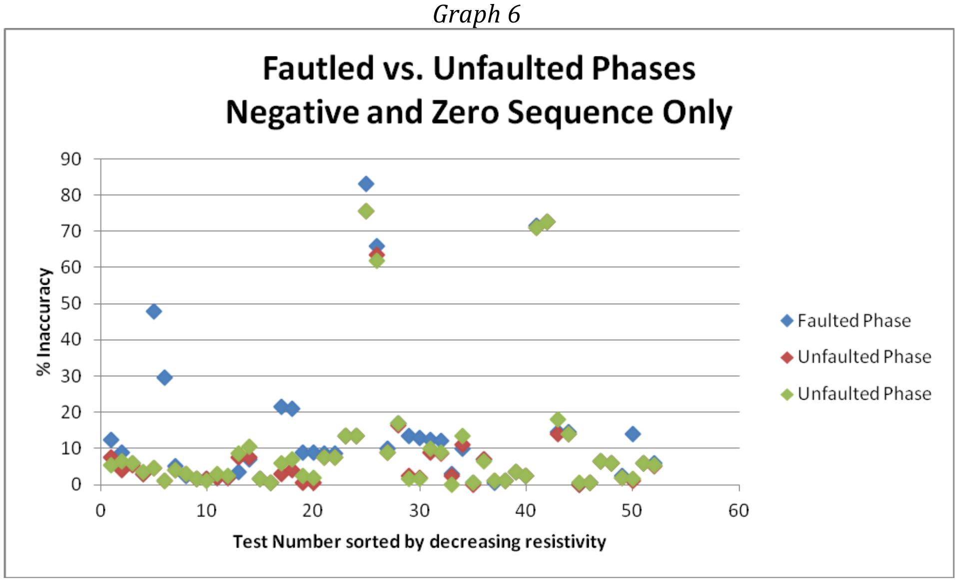

Graph 6A

Graph 6 shows accuracy broken down by using faulted versus unfaulted phases as your double ended reference phase. Results are sorted by decreasing resistivity. Utilizing the faulted phase as your refence phase appears initially to result in poor accuracy. However, when positive sequence tests are removed from the evaluation, as displayed in Graph 6A, the faulted phase performs comparably with zero and negative sequence, maybe only slightly worse. Even when including positive sequence tests, results only appear uniquely worse for resistive faults. Unfaulted phases, however, perform generally well across the board and should be used as best practice.

Fault Resistivity

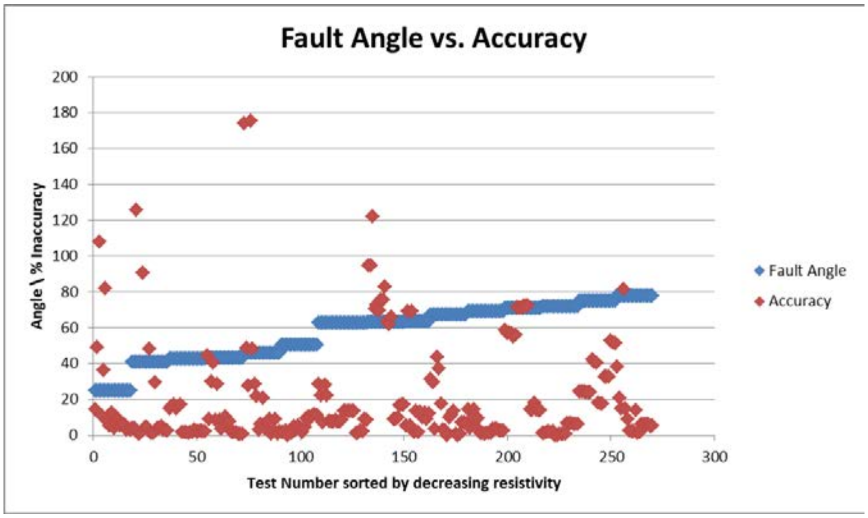

Graph 7

Graph 7 shows the relationship between Fault Angle and Accuracy. Fault Angle is inversely related to resistivity, so tighter fault angles can be interpreted as highly resistive faults. Fault resistivity appears to have a slight impact on location accuracy, especially in the very resistive region. Again though, as discussed in the faulted vs. unfaulted phases evaluation, positive sequence locations are prone to inaccuracy for resistive faults. With positive sequence tests removed, resistive faults have a minimal influence.

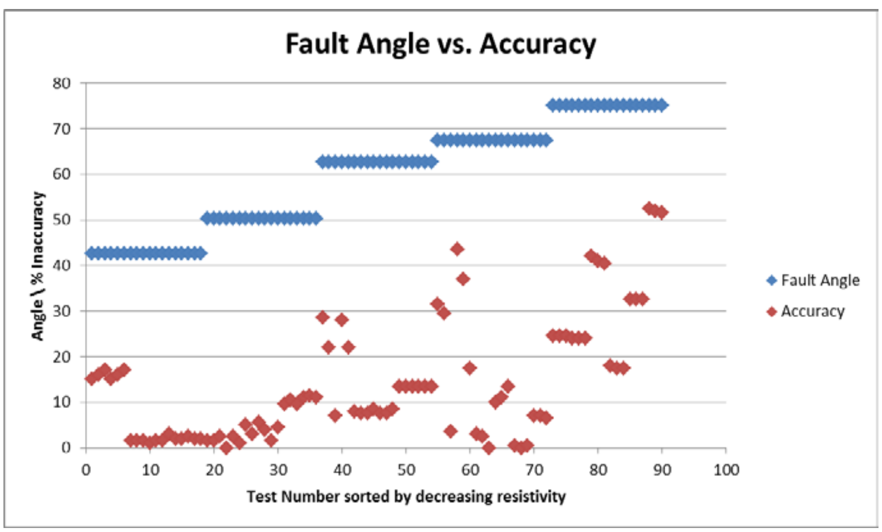

Graph 8

Interestingly, in our data set, 115kV faults showed a definite trend of decreasing accuracy with low impedance faults, as shown in Graph 8. High resistivity faults all came in with good accuracy, however, especially with positive sequence removed. This may be a consequence of peculiarities with our chosen faults rather than a general truth, however.

Outliers

Despite our best efforts to come up with a best practice for using the double ended algorithm, we still run into situations where the double ended method fails to provide a good result. One recent example involved a phase to ground fault located right at one terminal of the transmission line. The voltage on the faulted phase collapsed to less than 1% of nominal, which easily provides a close-in location to the trained eye. However this fault presented a challenge for the double ended method. Following our best practice, we utilized unfaulted phases as our reference for calculating the sequence components and looked at the results using positive, negative, and zero sequence values. Only one result using zero sequence values met our criteria, very little overlap and low result angles, but, the location given was 5 miles from the station on a 30 mile line. As mentioned previously, zero sequence is prone to little overlap, but tight angles generally imply accuracy. The fact alone that only one in six calculations even appeared accurate may imply all results are suspect. Perhaps zero voltage faults naturally confuse the equation, although it may well be something else entirely. To be sure, further research is needed in this area, as well as more examples of outliers like this.

Conclusions

- Performing the double ended calculation twice, reversing the near side and far side, provides additional data that aids in determining confidence.
- Precise time sync is not required as long as measurements are taken during the same fault state.
- Choosing a calculation window 1.75 cycles into the fault did not produce results that were statistically different than choosing a stable point in the fault.
- Referencing the unfaulted phase voltage for symmetrical component calculation provides better results, particularly when using positive sequence components.
- Zero sequence components were the most likely to be "in the ball park," within 20% accuracy.
- Negative sequence is the most likely to have high precision but is subject to notable inaccuracy.
- If negative sequence and zero sequence results largely agree, the negative sequence answer is the most likely to be highly accurate. This also gives a higher confidence that we have a good result.
- There is a very strong correlation between overlap and accuracy. In addition, visualization of overlap is very useful because it allows an analyst to spot it quickly. It should be noted, however, that zero sequence component calculations rarely produce overlap and may give a false sense of confidence in a result.
- There is a very strong correlation between low result angle and accuracy.
- Outliers occur that are difficult to predict.

Biographies

Patrick Hawks has worked at Dominion Virginia Power since 2008. He has a B.S. degree in Computer Engineering from Virginia Commonwealth University. His experience includes six years in system protection, wherein he participated in system modeling, relay settings, and protection standard development, implementing a substantial degree of automation, both in setting calculation and setfile generation. He has since moved into Fault Analysis and become involved in analyzing system events and web app development.

Amir Makki has worked at Softstuf since 1991. He has BS and MS degrees in Electrical Engineering from Tennessee Tech University and pursued his Ph.D. studies in Software Engineering at Temple University. His main interest is automating fault and disturbance data analysis. He is extensively published and holds a number of U.S. patents and trademarks. Amir is a senior member of IEEE and is an active member of the Protection Systems Relay Committee where he chaired a number of working groups including COMTRADE, COMNAME, and the Cyber Security Task Force for Protection Related Data Files.

Robert Orndorff has worked at Dominion Virginia Power since 1984. He earned an A.A.S degree in Electronics in 1986 and spent 11 years as a field relay technician and in 1997 transferred to the Fault Analysis department where he currently works. His current responsibilities include maintaining and configuring Dominion's Digital fault recorders, event retrieval and analysis from smart relays and DFRs. Robert is an IEEE member and has been a member of the Transient Recorder's User Council (TRUC) since 2002.

Maria Rothweiler has worked at Softstuf since 1991. She has a Bachelor of Arts degree in Computer Science from Temple University and pursued her associate degree in Mathematics at Bucks County Community College. Her main interest is developing software tools for display and analysis of power system fault and disturbance data. She is extensively published and holds a number of U.S. patents, trademarks, and copyrights. Maria is a member of IEEE and an active member of the IEEE Standards Association.

Brian Starling has worked at Dominion Virginia Power since 2003. He has a BS degree in Electrical Engineering from Virginia Commonwealth University. His experience includes testing, installing, and repairing relay systems and calculating line impedances and relay settings. He also has extensive experience analyzing and documenting transmission system operations. He is currently responsible for compliance reporting, misoperation investigations, and fault recorders. Brian is a Master Black Belt Six Sigma for helping reduced Dominion's transmission operations by 22%.

Kyle Thomas received his M.S. degree in Electrical Engineering from Virginia Tech in 2011 and is currently pursuing his Ph.D. while working for Dominion Virginia Power's Electric Transmission Operations Research group. He has technical expertise in power system protection/control, wide-area measurements, fault analysis, cascading analysis/physical security/resiliency, and system simulations. Kyle is a technical lead of Dominion's synchrophasor installations, applications, and training, and is actively involved in the North American Synchrophasor Initiative (NASPI), IEEE, and CIGRE organizations.