

B3/B5 Joint Session PS3 – Integration of Intelligence on Substations**Substation-based Waveform Analytics Monitoring System For Improved Circuit Awareness**

J. A. WISCHKAEMPER, C. L. BENNER, B.D. RUSSELL, K. MANIVANNAN
Texas A&M University
United States of America
jeffw@tamu.edu

SUMMARY

Increased use of sensors and sensing technologies in substations has resulted in ever-increasing volumes of data available to utilities. While many utilities believe this data contains useful information, extracting operational value from substation data is frequently challenging. In the case of waveform oscillography, which today can be obtained from a variety of instruments (relays, digital fault recorders, power quality monitors, etc.), transforming data into information traditionally requires significant effort by trained personnel with specific expertise in waveform analysis. Personnel frequently need to sort through hundreds or even thousands of oscillography records with no operational value to find only a handful with actionable information. Many utility companies may have information in their databases that would allow them to prevent faults and failures, but because of the skill required and effort involved, the existence of this information is usually discovered only after the event. Manual datamining is inefficient and time consuming, making real-time use impossible.

Researchers at Texas A&M have engaged in an ongoing, multi-decade monitoring project with over thirty utilities in five countries to capture and investigate failure signatures on distribution circuits. This database represents over 1,500 circuit years of exposure and has produced millions of high-fidelity oscillography records. Circuits are monitored on a one-per-feeder basis using conventional current and potential transformers. In close consultation with utility companies, researchers began to match electrical signatures to real-world circuit events, and quickly began to recognize and predict a variety of normal and abnormal events from their current and voltage signatures. Based on recorded signatures, it became clear that many faults and failures produce distinct, discernible signals for hours, days, or weeks before escalating into a more serious problem requiring utility intervention. While some failures cause high magnitude signatures, many incipient failures produce signals with very low magnitude, often smaller than normal circuit operations which are frequently missed by conventional data recorders. In many of these cases, alerting the utility to the existence of abnormal waveform signatures allows them to locate and correct the underlying issue before it caused a disruption to customers.

In the early stages of the project researchers relied on manual inspection and classification of waveform records to determine the cause of an event. As the project scaled, however, it became clear that it was not possible to manually examine each record, even for a team of several researchers and a relatively modest number of circuits. Because many incipient failure events produce relatively small perturbations in current and almost no change in voltage, triggering sensitivity must be set significantly lower than a traditional relay or fault recorder to extract the most operational value. This sensitive triggering results in nominally dozens of oscillography records on a typical circuit in a typical day. To cope with this volume of data, researchers created a generic waveform classification

and feature extraction engine which resides in substation data collection devices and automates the process of analysing waveform data.

The waveform classification engine takes generic waveform inputs and produces common language assessments of power system events (e.g. “motor start,” “capacitor on,” “breaker trip,” “probable failure of switch/clamp,” etc.). System operators and engineers have used these assessments to diagnose feeder events, enable condition-based maintenance, and evaluate the real-time health of circuits. On multiple occasions, the automated processing of waveform data has allowed utilities to improve operational outcomes by understanding when apparatus are beginning to fail or operate in an improper manner. This paper will explore both the waveform classification engine, as well as presenting multiple case studies where automated oscillography analysis was used to find and repair failing apparatus.

KEYWORDS

Waveform analysis, incipient fault detection, condition based maintenance, waveform monitoring

Introduction

Electrical apparatus on distribution circuits produce current and voltage transients and perturbations as they operate, both normally and abnormally. The nature of these transients is determined by the underlying physics of the system, and therefore provides information about the types of devices that are operating. For example, a motor generally produces a large inrush transient as it starts, but this transient has different characteristics than a load pickup transient after an outage is restored. A capacitor switching event produces high-frequency transients accompanied by a step change in voltage at the bus. Failure events also produce distinct and discernible characteristics, again based on the underlying physics of the system and failure mechanism. Like with signatures produced by normally operating devices, it is sometimes possible to use these signatures to infer the nature of the failure and its location on the system.

While some devices fail precipitously, other devices degrade and fail over an extended period – sometimes many hours, days, or weeks. In these scenarios, the electric system is still serving load to all customers, but careful observation of current and/or voltage waveform transients can identify that apparatus have begun to fail. In such cases, the opportunity exists for utility companies to find and fix apparatus before they cause an outage or other adverse event (e.g., fire, explosion, injury). For this to happen, the utility must 1) become aware that a device has entered an imminent failure mode, 2) be able to locate the device and, 3) act in a timely way to remedy the condition.

Until recently, awareness of developing incipient failure conditions was the primary limiting factor limiting improved outcomes with regard to true condition-based maintenance. In the United States, distribution circuits are primarily operated in a radial configuration, and may range from several line-kilometres to several hundred line-kilometres of exposure. An average distribution circuit contains thousands of components, ranging from connectors, fuses, insulators, transformers, conductors, poles, and the like. Most components have an expected lifetime of many decades, but all equipment eventually fails. Because of the large number of components and the fact they are spread across a wide geographical area, physical inspection of line apparatus is too costly and time intensive to be practical in most cases. Utilities might inspect certain high-value components (capacitor banks, load tap changers, etc.) on a periodic basis, but most equipment is operated in a run-to-failure mode: the apparatus is placed in service, operates for many years until it fails, at which point the utility replaces it. This statement is not intended as a criticism; as noted, until recently utilities have not possessed tools or resources to do anything better. In the past decade, however, utilities have deployed a variety of sensing technologies that produce waveform records during transient events on the power system. These data, sometimes referred to as waveform oscillography or point-on-wave data, vary widely in quality, but the goal is generally the same, namely to provide utility companies with additional information about the health and status of the system.

Historically, these waveform monitoring systems have generally fallen into two categories. One category of devices is broadly intended to be deployed on an as-needed basis when the utility suspects a problem may be imminent, usually at a particular location. Traditional on-site power quality (PQ) monitoring would be one example of this category. In this application, engineers typically install a device on a temporary basis – perhaps several days or weeks – then collect the device and analyse the data offline in the office for potential problems. A different category of devices are designed to be installed in an online fashion, but often feature waveform recording as an add-on to their core function. For example, many digital relays also produce oscillography records of faults. These records can be useful, but are secondary to the primary function of the relay, circuit protection. In the last decade, utilities have begun deployment of sensing technologies which are permanently installed and feature waveform recording as their primary function.

The advent of these systems has created a new problem, specifically data overload. When widely deployed across a utility's system, always-on sensing technologies can quickly produce thousands of waveform records per day. The vast majority of these records are innocuous and require no additional attention from the utility. Some, however, contain data that represent an opportunity for the utility to

improve outcomes compared to the traditional baseline, if that data can be operationalized. As such systems become more widespread, engineers are increasingly aware that collecting data is a necessary condition for improved situational awareness, but data collection alone is not sufficient.

Waveform analytics philosophy: finding the important needles in the needlestack

Over a two-year period, a particular distribution circuit experienced 178 individual fault events. As one might expect, many of these events were isolated in nature, and caused momentary protection operations which successfully reclosed with no sustained outage for customers. The utility which operated this circuit has communicating reclosers installed across their entire system and had an enterprise-level record of each of the 178 individual faults, the times they occurred, and what action, if any, was taken in response to the fault. What the utility *did not* know, based on its business-as-usual practices, was that 44 of the 178 faults occurred at the same location and had the same root cause.




Figure 1: Two identical faults 20 months apart

Figure 1 shows two of these faults, which occurred 20 months apart. In each case, the fault current magnitude, phase involvement, and recloser timing were substantially similar. The reason, as mentioned, was that vegetation repeatedly pushed the Phase A and B conductors together in a specific span when wind blew at sufficient velocity from a certain direction. The result was repeated customer interruptions, power quality problems, and eventually, a sustained outage after the line burned down.

The utility had installed an advanced waveform analytics system on this feeder but operated the system in a “blind trial” mode. That is to say, one group within the utility had access to the data and reports produced by the system, but no action was taken on these reports outside of the utility’s normal, business-as-usual practices. Technologically sophisticated utilities have increasing amounts of data available at their disposal, but seldom have tools to triage and operationalize that data. Such was the case here. The utility had records which could have alerted them to the incipient failure condition that eventually burned the line down, but personnel were unable to make sense of those records and identify which of the individual faults merited further attention. By contrast, the automated waveform analytics system immediately identified the incipient condition after the occurrence of the second fault and continued to report the problem for over two years following each of the 42 remaining occurrences, until the line burned down.

Research at over two dozen utility companies has shown that cases like this are particularly difficult for human operators to connect, absent automated tools. While some utilities have groups which

investigate fault records looking for patterns, the sporadic nature of these events and the fact they often occur weeks apart, frequently in storm conditions, interspersed with other unrelated faults makes this endeavour challenging at best. At a moderately sized utility with a few hundred circuits, it is not uncommon to have several thousands of momentary faults per year. Moreover, when one considers that conventional faults represent only a small portion of important events on a power system capable of causing incipient failures, the problem becomes more daunting still.

Consider another illustrative example. Many utilities employ switched line capacitors for voltage support. It is common for banks to switch using time based or volt-temperature control schemes, which generally result in a capacitor switching on sometime in the morning and off sometime in the evening. As noted in the introduction, capacitor banks produce a distinctive waveform transient when they switch on or off, which is relatively easy for humans to recognize, even with minimal training. Again, considering scale, a moderately sized utility with a few hundred circuits could easily have hundreds of capacitor switching operations per day during normal operating conditions – far too many for a human or group of humans to analyse manually. At the same time, switched line capacitors are known to have a very high failure rate – one utility with yearly inspections of distribution line capacitors cites that over 10% of banks have some failure which needs to be corrected each cycle. [1] A system which could automatically analyse capacitor switching operations for abnormalities has the possibility of both improving system operation and reducing maintenance costs for utilities.

Additional nuance is required when thinking about “normal” events like capacitor switching operations. For example, researchers have documented multiple failures of a capacitor controller where a bank begins to operate too frequently. In these cases, instead of one or two ON/OFF cycles per day, the controller may signal for the bank to cycle dozens, or even hundreds of times per day, prematurely wearing out switch contacts and creating power quality problems for all customers on the circuit and all circuits attached to the same bus. If examined from the level of single waveform recordings, each of these events are perfectly “normal” – the “abnormality” is the fact there are far too many of them. As a result, analytics systems need to recognize not only whether an individual event is “normal,” but also place that event in a larger context to know whether action is required.

Waveform Analytics Data Requirements

As noted in the introduction, the quality of waveform oscillography varies dramatically across devices. It should go without saying that if transients are not recorded, they cannot be analysed and therefore cannot provide value. Likewise, if a recording device captures a transient, but fails to do so with enough fidelity to accurately and distinctly characterize the conditions which would enable classification and location, the data records have little value. It is the authors’ experience, however, that many engineers simply assume that whatever devices they install will capture all important events and faithfully reproduce transients, without understanding the nuances and complexities involved. When engineers do examine specification sheets, they usually limit their queries to easy-to-quantify line items like sample rate and/or bits of resolution. While these quantities are important, they do not in themselves tell the entire story of whether a platform will be useful for automated waveform analytics. The concept of “sufficient fidelity” embodies a nexus of factors, all of which are important for an operational waveform analytics system. These include but are not limited to:

1. Always-on, 24/7/365 monitoring – some systems are intended to be deployed only on a temporary basis, or moved around from station to station. Other systems (e.g. line-powered sensors) only operate when there is sufficient current (generally tens of amperes). These systems are likely to miss incipient conditions either because they are not installed, or not powered when such events occur.
2. Near-realtime access to information – some systems record waveform data, but it is only retrieved or analysed on exception, when something is believed to have failed. Preventing incipient faults and failures requires data to be available before the catastrophic failure, not after it.
3. Sensitive triggering, especially on currents – many classes of incipient failures have precursors that may produce fewer than 5 amperes of line current, and almost no change in voltage. Few systems are designed to trigger sensitively enough to record such small

signatures. As stated previously, if signatures are not recorded, they cannot be used for analysis.

4. 19+ effective number of bits (ENOB) resolution – wide range and resolution are necessary to faithfully record high-magnitude faults (e.g., 10,000 amperes of primary current) while also faithfully recording the waveshape and characteristics of low-magnitude events.
5. Extended record lengths – many devices record relatively short records (i.e., cycles to tens-of-cycles). This is often sufficient for the analysis of simple, well-behaved faults, but frequently causes incorrect conclusions to be drawn regarding complex, evolving faults and incipient events, which, paradoxically, are the faults where detailed analysis is usually required.
6. Minimal “device fingerprinting” in recordings – Some devices, particularly those with lower sample rates, may include aggressive analogue filters for anti-aliasing. Depending on how these filters are implemented, they can cause substantial distortion to waveforms which can impact classification and location estimates. Additional caution should be noted regarding use of non-traditional sensing technologies.

Waveform Analytics in Deployment: Distribution Fault Anticipation

Beginning in the late 1990's, researchers at Texas A&M University began exploring how waveform recordings could be used to improve situational awareness for distribution utilities. Initial work occurred under the “Distribution Fault Anticipation” (DFA) project, funded by the Electric Power Research Institute (EPRI). The project began primarily as a data collection exercise to study naturally occurring failure signatures of line apparatus, and currently holds the largest known database of such signatures in the world, with over 1,500 circuit years of real-world exposure at over three dozen utility companies in five countries. Distribution circuits were monitored on a per-feeder basis with specialized, high-fidelity waveform monitors using traditional bus potential transformers (PTs/VTs) and feeder current transformers (CTs). In early stages of the project, researchers worked in close collaboration with utility companies to match electrical signatures to real-world events. After some time, researchers found they were able to recognize and predict certain events based on their current and voltage signatures.

Even in the early stages of the process, it became clear that manual inspection of each waveform record with the intent to classify it would not be practical at scale. As a result, the project quickly shifted toward developing software techniques which could automatically triage “normal” events like motor starts and load transients, and help researchers focus on events that were likely to be more consequential. After collecting statistically significant numbers of individual events for various types, researchers began to analyse features unique to specific events with the intent of designing algorithms which could classify disturbances with high accuracy. A full discussion of algorithms and algorithm development is outside the scope of this paper. [2] This process resulted in the development of a “black-box” classification engine which accepts as input waveform recordings and outputs “human readable” information about the event contained in the waveform. In cases where abnormal operations are detected, the classification engine also displays parameters which can be useful for helping system operators locate the underlying condition.

The ultimate goal of the DFA waveform analytics engine is to reduce the amount of background knowledge required to make effective use of recorded waveform data. Said differently, the goal is to make the system useful to an operator even if they do not have any specialized training in waveform analysis, as is frequently the case for utility users, especially those in control room settings. While no system can replace human judgment, reducing the amount of technical experience required to extract value from waveform data is critical for real-world success.

Waveform Analytics: Real-world Examples

The following recent case studies illustrate how waveform analytics can be used to improve situational awareness of circuit events. All cases were taken from operational circuits which utilized Distribution Fault Anticipation devices for waveform recording and automated analysis.

Case Study 1: Animal caused repetitive fault.

During a period of normal system operations, a circuit experienced a fault with unusual characteristics, show in Figure 2. Researchers noted the unusual nature of the fault – namely that it began as a Phase A-to-ground fault, but evolved into an AB-ground fault after approximately 10 cycles. This fault tripped the substation breaker, interrupting power to over 1,200 customers before successfully reclosing.




Figure 2: Initial fault, Case Study 1

During normal conditions there would be no need to take further action based on this information, and indeed the utility did nothing. Faults which transition from one phase to involve multiple phases are unusual, but not uncommon enough to merit further investigation. Six days later, however, a second fault occurred with substantially similar characteristics, shown in Figure 3. This fault was automatically flagged by the waveform analytics engine as a possible repetitive fault, suggesting that additional action was required.




Figure 3: Second fault, Case Study 1

At this point the waveform analytics system had recorded both events, classified them individually as AB-ground faults, and noted that the events were likely related to each other. It automatically

generated a report, shown in Figure 4, which allowed an engineer to dispatch a crew for further investigation.

Breaker recurrent fault

AB

Breaker fault, 2178 Amps

2 (6 days)

2021-11-08 09:07:34

Export

Show

10

entries

Search:

Event Type	Phases	Comments	Count	Last Occurred	Graph
Breaker reclose	AB	F-(71.5c,2178A,ABN)-T-40c-C	1 op	2021-11-08 09:07:34	
Breaker reclose	AB	F-(70.5c,2224A,ABN)-T-40c-C	1 op	2021-11-02 09:12:33	

Page #

First

<

1

>

Last

Figure 4: Automated report generated by waveform analytics engine

Locating the event was a relatively straightforward process. The report in Figure 4 contained a fault current estimate for a LLG fault of approximately 2,200 amperes. This information alone, combined with a competent system model and engineering analysis software would allow an operator to dispatch a crew to at most a handful of locations. A closer look at the waveforms, however, provided additional clues which helped guide the search. First, researchers noted that the relatively swift transition of both faults from LG->LLG would require a spot on the system where at least two phases and ground were physically close together, such as in a transformer, or at a recloser. Additionally, careful inspection of the second event showed that the fault tripped both the substation breaker and a midpoint recloser, suggesting the location of the fault was likely near a recloser, or in a zone where the two schemes were close to overlapping. Based on the information provided, the utility initiated a search, and found the cause of the fault within hours of the second occurrence. The responding crew found a dead squirrel fused to an energized recloser case, which appeared to have been dead approximately a week, shown in Figure 5. Prompt action to locate and remove the squirrel likely prevented future faults and interruptions to all customers on the circuit, as well as damage to the recloser and surrounding apparatus.



Figure 5: Dead squirrel on recloser

Case Study 2: Capacitor arcing

Arcing in the path of a capacitor bank is a particular and problematic form of “series arcing.” Series arcing occurs when the contact surfaces of a current carrying device develop an abnormally high impedance, which results in concentrated heat and surface-to-surface arcing as load current passes through the device. Unlike traditional “shunt” arcing, there is no unintentional current flow. Some series arcing events can progress from initial electrical indications to final device failure in tens of minutes to tens of minutes, but many failures take place over days or even weeks. During that time,

load continues to be served, although with intermittent and heavy distortion of the supply voltage downstream of the failing device.

When series arcing occurs in the path of a capacitor switch, it roughly approximates opening and closing a switch every half cycle and can produce substantial voltage transients. These transients are seen by all devices on the bus and have been observed to propagate even to other substations over the transmission system. Capacitor arcing has precipitated the failure of other line devices including lightning arresters and capacitor banks, both on the same circuit as the failing capacitor, as well as other circuits served from the same bus.

Early one morning, a three-phase, 900-kvar switched capacitor switched off. Two seconds after the successful switching operation, Phase B began conducting again, but with arcing. This signature is consistent with a partial loss of vacuum in a vacuum switch attached to a capacitor bank.




Figure 6: RMS and oscillography, capacitor arcing

Figure 6 shows both RMS and waveform plots for the capacitor arcing event. Unlike some capacitor arcing events, this particular instance is remarkable for how little the current and especially voltage were affected, even during periods of substantial arcing. Specialized signal processing in the waveform analytics engine reliably identified these signatures as capacitor arcing and presented them on the utility dashboard. Alerted to the developing problem, the utility dispatched a crew the same day, which found the damaged capacitor switch shown in Figure 7 and repaired the switch.




Figure 7: Damaged capacitor switch

Conclusion

Automated waveform analytics have the potential to revolutionize how utilities operate their systems by vastly improving situational awareness. Because voltage and current waveforms contain information about the real-time health of the system, they provide utilities opportunities to fix equipment in many cases before it reaches a final failure, rather than afterward.

To extract maximum value, waveform analytics systems must:

1. Capture most to all events of interest, which requires triggering an order of magnitude more sensitive than most devices available on the market. **If the data is not recorded, it cannot be analysed or operationalized.**
2. Provide access to the data in near-real time. **If data is not analysed and communicated before the failure, it is of no operational value.**
3. Automate the data analysis pipeline such that personnel do not need to look at every waveform record. **If manual analysis is required, the system will quickly become unscalable. Records which are not analysed cannot be operationalized.**
4. Reduce the amount of knowledge required for users to extract value from the system. **The system should not require users to have extensive training in waveform analysis to extract meaningful value.**

References

[1] Final Report, EPIC Project 2.34 – Predictive Risk Identification with Radio Frequency (RF) Added to Line Sensors. Pacific Gas and Electric Company, 11-12-2020. Retrieved at (https://www.pge.com/pge_global/common/pdfs/about-pge/environment/what-we-are-doing/electric-program-investment-charge/PGE-EPIC-Project-2.34.pdf)

[2] Manivannan, K. “Power Distribution System Event Classification Using Fuzzy Logic.” Texas A&M University Dissertation, 2012.