

Distribution Fault Anticipator

Phase II, Algorithm Development and Second-Year Data Collection

Technical Report

Distribution Fault Anticipator

Phase II, Algorithm Development and Second-Year Data Collection

1001688

Interim Report, February 2005

EPRI Project Manager A. Sundaram

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATION(S) NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

ORGANIZATION(S) THAT PREPARED THIS DOCUMENT

Texas A&M University

ORDERING INFORMATION

Requests for copies of this report should be directed to EPRI Orders and Conferences, 1355 Willow Way, Suite 278, Concord, CA 94520, (800) 313-3774, press 2 or internally x5379, (925) 609-9169, (925) 609-1310 (fax).

Electric Power Research Institute and EPRI are registered service marks of the Electric Power Research Institute, Inc.

Copyright © 2005 Electric Power Research Institute, Inc. All rights reserved.

CITATIONS

This report was prepared by

Texas Engineering Experiment Station Department of Electrical Engineering Texas A&M University College Station, TX 77843-3128

Principal Investigators C. Benner K. Butler-Purry B. Russell

This report describes research sponsored by EPRI.

The report is a corporate document that should be cited in the literature in the following manner:

Distribution Fault Anticipator: Phase II, Algorithm Development and Second-Year Data Collection, EPRI, Palo Alto, CA: 2005. 1001688.

PRODUCT DESCRIPTION

As part of a continuing project to use measured electrical signals to determine when power system equipment is deteriorating, eleven utilities monitored feeder lines using a prototype Distribution Fault Anticipator (DFA). This interim report describes the second year of data collection and the development of algorithms to automate the analysis of collected data.

Results & Findings

Because the electrical changes that deteriorating equipment produces are often very subtle, a large number of normal as well as abnormal system events must be analyzed in order to ensure the robustness of recognition methods. Under the present project, the research team designed a new prototype platform and delivered fourteen prototypes to eleven utility companies for installation in substations in various locations in the continental United States and Canada. Each prototype is capable of monitoring from one to eight feeders in a given substation. Participating utilities ordered prototypes configured with the number of feeders they desired, based upon the substations in which they planned to install the prototypes. The total number of feeders monitored under the present project is between sixty and seventy. The larger number of prototypes to a wide variety of operating philosophies and practices. In addition to operational diversity, geographical diversity provides further variations in the operating conditions. To handle the flood of data from the prototype system, algorithms are under development to classify events, focusing initially on normal system events to provide a baseline for the identification of significant anomalies.

Challenges & Objectives

Earlier EPRI-sponsored efforts at Texas A&M University proved the concept of using sensitive monitoring to provide early warnings of impending failures (EPRI report 1001879) and began the process of involving utilities, substations, and feeders in the further development of the technology (EPRI report 1002153). The present project builds on this work by continuing to expand the database of monitored events and by developing algorithms to automate the analysis process to make fault anticipation useful in day-to-day operations.

Applications, Values & Use

Distribution Fault Anticipator technology offers the potential to operate and maintain distribution circuits in a new and exciting way. Imagine having the ability to know that feeder equipment is deteriorating and will disrupt service in the near future. A tool that could provide this type of information would enable fundamental changes in operating and maintenance philosophies. Certain types of expensive, low-return activities that used to be necessary in order to maintain an

adequate level of service quality and reliability could be reduced or eliminated. This tool has the potential to allow utilities to improve reliability and reduce O&M cost at the same time.

EPRI Perspective

Competition drives us all. Today's electric utility company must do things better and more efficiently than the competition. Historical practices that kept the lights on still may be effective, but often are too inefficient and costly for today's market. In addition, reliability indices are becoming increasingly important drivers in the competitive environment. Projects like the Distribution Fault Anticipator have the potential to improve reliability by minimizing outages that would have occurred due to equipment that was going to fail.

Approach

Building on previous work, researchers designed a more robust prototype hardware and software platform for distribution fault anticipation. Eleven utilities are participating by installing these Prototypes on their circuits for a nominal two-year period and by serving as the hands-on operators of the Prototypes. The extensive monitoring program will provide a database of measurements for a wider variety of failure modes, complete with detailed documentation by the utility participants. In addition, the research team is devising methods to automate data analysis in order to convert measured data into information useful to operating utilities, thereby making the technology more practicable. The ultimate goal is not to provide the utility with a mountain of data, but rather to automatically sort through that mountain to provide useful information about the health of the system.

Keywords

Reliability Equipment failures Fault anticipation Incipient faults Condition-based maintenance

ACKNOWLEDGMENTS

This project enjoys the active participation and support of eleven utility companies. Each of these companies has committed time and resources to acquire and install equipment, allocate significant levels of personnel efforts, and provide other support and guidance along the way. This support is key to the success of an undertaking such as this and both Texas A&M University and the Electric Power Research Institute gratefully acknowledge these companies and their contributions. Utility companies participating in this way are:

- Southern Company/Alabama Power Company
- BC Hydro
- City Public Service, San Antonio, Texas
- Exelon
- Consolidate Edison
- Keyspan Energy
- MidAmerican Energy
- Northeast Utilities
- Omaha Public Power District
- TXU Electric Delivery
- Tennessee Valley Authority/Pickwick Electric Cooperative

CONTENTS

1 INTRODUCTION	1-1
Historical Market Perspective	1-1
Project Background	1-2
Reference	1-3
2 GOALS OF ONGOING WORK	2-1
Expand Database of Incipient Events	2-1
Enhance Quality of Database of Incipient Events	2-3
Automate Data Management	2-3
Automate Data Analysis	2-4
3 WORK TO DATE	3-1
DFA Prototype Design and Production	3-1
Design Overview	3-1
DFA Prototype Field Unit	3-1
DFA Prototype System Data Storage	3-3
DFA Prototype Master Station	3-3
Initial Problems Encountered and Corrected	3-5
Users Group Meetings	3-6
Automatic Classification Algorithms and System	3-8
Classification Algorithms	3-8
Automated Classification System	3-10
4 FIELD RESULTS	4-1
Data Collection Philosophy	4-1
Summary of Failures Documented to Date	4-2
Case Studies	4-2

Malfunctioning Capacitor Bank Controller	
Repetitive Overcurrent Fault	
Long-Term, Repetitive Arcing	4-11
Failing Switches	4-14
Short-Lived, High-Current Events	
Deceptive Capacitor Fault	
Arcing on Secondary Customer Service Cable	4-19
Sticking Load Tap Changer (LTC) Controller	
Inoperative Capacitor Bank Switch	
Temporarily Stuck Capacitor Bank Switch	
Intermittent Capacitor Switch Restrike	
5 PRESENT STATUS AND FUTURE WORK	5-1

LIST OF FIGURES

Figure 3-1 DFA Prototype Field Unit Functional Block Diagram	3-2
Figure 3-2 Data Capture Viewed in DFA Prototype Master Station GUI	3-4
Figure 3-3 Statistical Data Viewed in DFA Prototype Master Station GUI	3-5
Figure 4-1 Phase Capacitor Short Circuit	4-4
Figure 4-2 Continuity Failure in Capacitor Oil Switch Contacts	4-5
Figure 4-3 "Continuous Transients" During Capacitor Switch Failure	4-6
Figure 4-4 Final Failure of Capacitor Bank, After Four Days of Intermittent Connection	4-7
Figure 4-5 Second in Series of Overcurrent Faults from Tree Limb Contact	4-8
Figure 4-6 Final in Series of Overcurrent Faults from Tree Limb Contact	4-9
Figure 4-7 Tree Limb That Caused Repetitive Faults and Burned Line Down	4-10
Figure 4-8 Early Episode of Repetitive, Intermittent Arcing	4-11
Figure 4-9 Episode of Repetitive, Intermittent Arcing, Almost a Month Later	4-12
Figure 4-10 Measured Signals Produced by Failure of Cutout	4-14
Figure 4-11 First Instance of High-current, Short-duration Fault	4-16
Figure 4-12 Second Instance of High-current, Short-duration Fault	4-16
Figure 4-13 First Instance of Capacitor Fault	4-18
Figure 4-14 Second Instance of Fault of Same Capacitor	4-18
Figure 4-15 Arcing Measured on Secondary Service Cable	4-20
Figure 4-16 High-frequency Current from Arcing on Secondary Service Cable	4-20
Figure 4-17 Load Tap Changer Stepping Excessively	4-22
Figure 4-18 VARs Recorded When Balanced Capacitor Bank Switched ON	4-24
Figure 4-19 VARs Recorded When Unbalanced Capacitor Bank Switched ON	4-24
Figure 4-20 RMS Amps Recorded When Unbalanced Capacitor Bank Switched ON	4-25
Figure 4-21 VAR Flows Caused by Intermittently Sticking Switch	4-26
Figure 4-22 Restrike Current When Capacitor Bank Switched OFF	4-27

1 INTRODUCTION

Historical Market Perspective

The electric utility of today is not what it was ten years ago. The coming decade will accelerate change further, as deregulation and competition become the norm. Electric utilities must adapt to the changing marketplace in order to survive and prosper.

The historical position of a utility as a regulated geographic monopoly offered utilities and their investors steady, predictable rates of return on their investments. As long as the utility provided an acceptable level of service by doing things they always had done, there was little incentive to seek innovative avenues of reducing operating costs.

Utility companies perform maintenance on periodic schedules that they have developed over long periods of time. Analysis of the historical performance of certain types of equipment provided statistical information about "safe" maintenance intervals. By observing these intervals, utilities reduced the likelihood of certain types of potentially catastrophic failures. Therefore, when determining maintenance intervals, the primary focus was on maintaining service, with only limited attention to cost. Obviously, performing maintenance on all equipment of a given type at an interval short enough to ensure an acceptably low rate of failure means that much of the equipment will be maintained more often than necessary.

Today, there are ever-increasing demands to lower electricity rates in order to attract and retain customers. At the same time, utility companies must operate profitably or go out of business. An obvious solution is to reduce expenses, both capital and operating. Because of its sheer size, the operating and maintenance (O&M) budget is a tempting target.

Breakers are maintained for a reason. Trees are trimmed for a reason. Patrols are conducted for a reason. When evaluating the reduction or even elimination of certain maintenance activities and expenditures, one must consider that each of these activities has a purpose and that arbitrarily reducing or eliminating it will affect the ongoing reliability of the system.

The utility company faces a conundrum: How do I target limited resources so that I significantly reduce O&M expenses while also identifying critical problems on my circuits before they cause service disruptions? This project seeks a solution to just that problem.

Project Background

EPRI and Texas A&M University's (TAMU's) Power System Automation Laboratory have a long and productive history in the area of applying advanced technologies to challenging power system problems. Much of this work focused on the development of techniques for the detection and recognition of high-impedance arcing faults.

During the course of their work on high-impedance fault detection, researchers noted numerous instances of changes in certain electrical parameters that were followed some time later by failures on the power system. For example, they observed instances of signals indicative of arcing that preceded fair-weather faults on the system hours or days later. As another example, they noted a correlation between increased arcing signals and insulator contamination levels.

Based in part upon this work, TAMU and EPRI postulated that it might be possible to detect and use changes in electrical parameters to "anticipate" when failures and faults were likely to occur. They conducted initial proof-of-concept research in which they confirmed that various types of failing equipment produce measurable changes to various electrical parameters. They recorded instances of electrical changes indicative of problems developing in several types of equipment, including a failing lightning arrestor, a failing substation voltage regulator and a chattering capacitor bank switch. They also measured distinct parametric changes that occurred coincident with right-of-way tree trimming on distribution circuits. In addition, they were able to recognize intermittent overcurrents that caused numerous voltage sags and momentary interruptions to downstream customers. [1]

In short, the proof-of-concept phase successfully demonstrated the potential to measure electrical changes that accompany deteriorating distribution system equipment. During the process, the research team developed considerable expertise in the diagnosis of measured events. They noted that most events on the monitored feeders were normal events (e.g., capacitor banks switching, load tap changers stepping, motors starting, etc.). Diagnosis of large numbers of such events is tedious and time consuming, but critical to the identification of the relatively small number of events of real interest, i.e., those that represent incipient faults.

In summary, the conclusions of the proof-of-concept phase were as follows:

- Deteriorating or contaminated equipment often produces electrical changes.
- These electrical changes often are measurable and can be used to recognize the deteriorating conditions.
- Anticipating faults has the potential to allow utility companies to reduce expenses while at the same time maintaining or enhancing service quality and reliability.
- The electrical changes that deteriorating equipment produces often are very subtle and require high-fidelity conditioning and signal processing for their measurement and analysis.
- Normal system events cause changes to some of the same parameters to which incipient faults cause changes, but detailed analysis of these changes provides the basis for differentiating them.

• Manual processing of the data is tedious and time consuming, making it impractical to use this process on a widespread basis. However, it should be possible to automate much of the process through the proper use of technologies such as signal processing, pattern recognition, artificial intelligence, etc.

Based upon these observations and conclusions, TAMU and EPRI began the current project. Subsequent chapters in this interim report provide more detail about the methodologies employed, results from field installations to date, and the status of developing algorithms and methods for recognizing events automatically.

Reference

1. Distribution Fault Anticipator, Phase I Final Report, EPRI, Palo Alto, CA: 2001. 1001879.

2 GOALS OF ONGOING WORK

The overall goal of this project is to develop the use of measured electrical signals more fully, to determine when power system equipment is deteriorating. This would enable utility companies to anticipate faults and take corrective actions to prevent adverse consequences to the system and to their customers. Ultimately, this has the potential to enhance service quality and reliability, while reducing the need for routine maintenance activities and the large expense associated with them. The following subsections detail the project's main tasks in light of this goal.

Expand Database of Incipient Events

During the proof-of-concept phase ("Phase I") of the project, the research team designed and constructed a small number of pre-prototype data collection devices. Where possible, they used off-the-shelf components in order to minimize design time and cost, with relatively little attention given to the per-unit cost of the equipment. They built and installed four pre-prototypes at utility companies across the United States. Each pre-prototype monitored one feeder.

Over a period of approximately two years, these four pre-prototypes recorded the following types of events:

- Normal system events (capacitors switching, load tap changers operating, etc.)
- Conventional overcurrent faults
- Lightning arrestor deteriorating and finally failing catastrophically
- Capacitor bank switch "bouncing" or "chattering"
- Substation voltage regulator failing

In addition, the four pre-prototypes recorded two examples in which intermittent short circuits caused repetitive overcurrent faults. One case was diagnosed as an intermittent contact between the overhead primary and a ground point; the other remains unresolved. Finally, they measured several instances of parametric changes that accompanied routine right-of-way tree trimming.

These measured events provided a vivid demonstration of the types of events that cause electrical changes that could enable fault anticipation. However, the database of events recorded in Phase I does not, in general, contain a multiplicity of instances of the different types of failures. Specifically, the database contains one instance of a failing lightning arrestor, one instance of a bouncing switch and one instance of a failing voltage regulator. Obviously, it would not be prudent to base any diagnostic system on such a small sample of recorded events. A central focus

Goals of Ongoing Work

of the present project, therefore, is to expand the database of recorded events significantly, both in number and in variety.

Under the present project, the research team designed a new Prototype platform and delivered fourteen Prototypes to eleven utility companies for installation in substations in various locations in the continental United States and Canada. Each Prototype is capable of monitoring from one to eight feeders in a given substation. Participating utilities ordered Prototypes configured with the number of feeders they desired, based upon the substations in which they planned to install the Prototypes. The total number of feeders monitored under the present project is between sixty and seventy. For various reasons internal to various utility companies, installation of the Prototypes proceeded at various rates, with several installed in early 2003, while the last unit was not installed until late 2004. In addition, being a research-grade platform, there were several hardware problems that affected the Prototypes, limiting the time in service for some of the units during the early months of the project.

The larger number of Prototypes and monitored feeders is increasing the variety and number of events in the database. The diverse set of participating utilities exposes the Prototypes to a wide variety of operating philosophies and practices. In addition to operational diversity, geographical diversity provides further variations in the operating conditions. All of this is good for development of robust analysis and classification methods.

This project is expanding the variety of recorded events in the database in at least three ways:

- <u>**Types of equipment**</u> Phase I recorded instances of deterioration and failure for several types of equipment, but certainly not for all types of equipment. For example, the original database contained measurements from a failing lightning arrestor, but not from a failing secondary service cable. The present effort has added this latter type of event, and others, to the database.
- <u>Failure mechanisms</u> A given type of equipment may have multiple failure mechanisms. For example, a porcelain insulator might fail because of mechanical damage (e.g., from natural causes or vandalism). However, it also might fail because of contaminant buildup (e.g., salt spray, sand, or chemicals). The expanded number of monitored feeders increases the likelihood of measuring signals from multiple failure mechanisms for given types of equipment. This is important because different failure mechanisms may affect monitored parameters in different ways and, thus, may require different recognition techniques.
- **Quantity of events** As previously mentioned, the original database contained several types of equipment failure, but in general, not multiplicities of given types of failures. The larger number of monitored feeders increases the likelihood of experiencing multiple instances of specific types of equipment failures and their underlying failure mechanisms. In addition, it is important to experience a large number of normal system events, in differing operating environments, in order to ensure the robustness of recognition methods. This is important in the design of automated recognition methods, so that the methods can be robust enough to ignore minor variations between power systems and between individual pieces of equipment.

Enhance Quality of Database of Incipient Events

During Phase I, participating utilities installed pre-prototype systems in their substations. They also provided information about events on their systems when the research team at Texas A&M made specific inquiries. However, they had no day-to-day responsibility for formally investigating and documenting events on their systems.

In the current project, participating utility engineers bear primary responsibility for day-to-day operation. Each utility is responsible for allocating a significant portion of an engineer's time and effort for the duration of this project. The engineer is responsible for routinely checking the Prototype for newly recorded events and for investigating and determining the cause of each event. The engineer then is responsible for documenting the underlying cause of the event and for adding this information to the database. Obviously, having large numbers of recorded events in the database is of limited value if the relatively few events of real interest do not have thorough identification and documentation.

Automate Data Management

During Phase I, the research team periodically (e.g., several times per week) initiated contact with each of the pre-prototypes from a master computer at Texas A&M headquarters. They did this via dial-up modem connections. They then retrieved new data from the pre-prototypes. This process was tedious and time consuming, but, for the limited number of pre-prototypes in the proof-of-concept phase, it was satisfactory.

With more than a dozen Prototype systems in the field recording large volumes of data from over sixty feeders, the amount of human intervention necessary to retrieve field data must be kept to a minimum in the current project. In addition, whereas TAMU was the sole repository for collected data during Phase I, utility engineers need access to the data from their Prototypes in order to perform their analysis, investigation, and documentation responsibilities.

The Prototypes utilize the following technologies to automate data collection:

- **<u>Reliable high-speed communications</u>** DFA Prototype Master Stations at TAMU and at utility engineers' offices communicate with the corresponding Field Units via high-speed Internet, utilizing industry-standard TCP/IP Sockets communications.
- <u>Automated polling</u> Each DFA Prototype Master Station is configured to poll the DFA Prototype Field Units periodically and to collect from them any information recorded since the last polling cycle. Of necessity, this occurs with no human intervention.
- <u>Integration of classification results with data captures</u> The large number of recorded events, most of which represent normal system switching events, makes it important to have means to integrate the results of manual and automatic classifications with the data records themselves. The DFA Prototype Master Station software allows specific classification information to be tied to each recorded event. Further, during automated polling cycles, the results of utility classification efforts are shared with TAMU, and vice versa.

Automate Data Analysis

During Phase I, the research team's emphasis was on demonstrating the concept of recognizing signatures of failing equipment. The team became proficient at recognizing a large number of normal system events and certain characteristics of incipient failures. The bulk of their analysis consisted of visual examination of waveforms and of extracted parameters in a custom data-viewing program that they developed. This provided them with the ability to view the behavior of a large variety of data over time, including current and voltage waveforms, computed power values, and harmonic and non-harmonic frequency components. The team performed more detailed analyses from time to time, but they generally used generic third-party tools (e.g., spreadsheets, MATLABTM, etc.) for these purposes. They did not attempt to automate the analysis techniques.

The manual analysis approach was appropriate for the proof-of-concept goals of Phase I. However, it was time consuming and required the development of a high level of expertise, making it unfeasible for widespread application by utility companies. The amount of time to develop the expertise would be considerable and, more importantly, the amount of time needed to perform the analyses on a day-to-day basis would be prohibitive on an ongoing basis.

For these reasons, one of the central goals of the current project is to automate as much of the analysis process as possible. The research team has begun to develop algorithms for automatically classifying events. The early months of this effort have concentrated to a large extent on identifying normal system events, such as capacitor banks switching, load tap changers operating, motors coming on, etc. There are two practical reasons for this approach, both of them having to do with the fact that there are so many of these normal events compared to the relatively small number of abnormal events. First, one of the other goals of the project is to gather sufficient data to develop robust recognition techniques. Trying to develop these techniques before the supporting data are collected would be putting the horse before the cart. Second, the large number of normal events that participating utility engineers must classify each day can be daunting. Automated recognition of these events eases this burden, at least to some extent.

3 WORK TO DATE

Work to date has consisted of three main tasks: producing and deploying Prototypes at host utility substations, assisting utility engineers in the initial analysis of collected events, and developing automated analysis routines. The following paragraphs discuss these activities in more detail.

DFA Prototype Design and Production

Design Overview

The initial task of the project was the design and production of the Prototype platform. Where practical in terms of both time and cost, the design team used commercially available, off-the-shelf components as building blocks. The reason was that development of the data collection platform was a means to an end, not the end in itself. With this in mind, however, certain specialized components were not available as off-the-shelf components. In these cases, the design team did the required low-level design and production.

The following subsections describe the two components of each DFA Prototype system: the DFA Prototype Field Unit and the DFA Prototype Master Station. In order to avoid being overly verbose, here and elsewhere in this document, the term DFA Prototype Field Unit often is shortened to Field Unit and the term DFA Prototype Master Station often is shortened to Master Station.

DFA Prototype Field Unit

Each DFA Prototype Field Unit ("Field Unit") resides in a 19" rack-mount chassis with nine card slots. One slot is reserved for a Master CPU card, which consists of an off-the-shelf single-board computer (SBC) with onboard RAM (random access memory) and a hard disk for storage of configuration information and long-term data. In addition, the Master CPU card provides for communication with an off-the-shelf weather station that is a peripheral to each Field Unit. Other peripherals include commercial-grade monitor, keyboard, mouse, and UPS (Uninterruptible Power Supply).

Each of the other eight slots is available for insertion of a feeder module. Each feeder module is designed to monitor the current and voltage signals for one feeder. Each feeder module provides four current inputs (five-amp nominal) and three voltage inputs (120-volt nominal), for connection to the secondary windings of conventional current and potential transformers (CTs

Work to Date

and PTs). Each feeder module consists of two main components. The first is a custom-designed signal-conditioning module that interfaces with the CT and PT inputs, provides appropriate signal conditioning, and converts the input signals to digital format at a rate of 15,360 samples per channel per second (i.e., 256 samples per cycle at 60 Hertz). The second is a single-board computer (SBC) with onboard RAM. The feeder module's SBC communicates with its signal-conditioning module via PC-104 bus. This allows the feeder module's SBC to adjust various settings on the signal-conditioning module and allows the signal-conditioning module to send its digitized samples to the feeder module's SBC. The feeder module's SBC performs various numerical functions on the data, including computation of standard power system quantities, frequency components, etc. It also monitors the levels of various parameters and triggers high-speed data captures when real-time values cross configurable thresholds. Finally, the feeder module's SBC is responsible for calculating various statistical values that it periodically records to long-term storage.

The main functions of the Master CPU card are to coordinate the activities of the multiple feeder modules, manage long-term data storage, and communicate with the outside world. The feeder module CPUs communicate with the Master CPU via an Ethernet hub inside the Field Unit's chassis. When a feeder module needs to store information in long-term storage, it notifies the Master CPU of this and then sends the data to the Master CPU, which in turn carries out the actual storage operation.

The other component of the Field Unit is a weather station. The weather station is a commercially available, off-the-shelf component. It provides a suite of weather sensors that connect to a weather station console that resides inside the control house. This console provides a serial connection that allows the Field Unit's Master CPU to poll it for real-time weather values. The measured weather parameters include temperature, relative humidity, wind speed and direction, and rainfall.

Figure 3-1 illustrates the interconnection of the various components of each DFA Prototype Field Unit.



Figure 3-1 DFA Prototype Field Unit Functional Block Diagram

DFA Prototype System Data Storage

Each Field Unit continuously conditions, digitizes, and analyzes multiple channels of incoming data from each of its feeder modules. Obviously, it is not practicable to record all incoming data on a continuous basis. Therefore, the Field Units store two types of data:

- 1. Statistical data The Field Unit continuously calculates numerous parameters from the incoming voltage and current signals. Depending on the particular parameter, it calculates these values at intervals ranging from a fraction of a cycle to multiple cycles. At the end of each configurable interval, nominally 15 minutes in length, each feeder module CPU calculates several statistics from each of these parametric data streams. The Field Unit then stores each feeder module's statistical values in its database, for long-term storage, retrieval, and analysis.
- 2. Data captures In addition to time-driven statistical data, the Field Units also store highspeed waveform data from time to time. Each feeder module determines the need to store this high-speed data when any one of the monitored parameters crosses one of its configurable thresholds. When this occurs, the feeder module records high-speed waveforms to long-term storage for several seconds.

DFA Prototype Master Station

By design, each Field Unit stores a significant amount of information each day. The purpose of the Master Station is to retrieve this information and to allow the user a straightforward way to analyze and manage it. Conceptually, the Master Station is the user's window to the Field Units.

The Master Station hardware consists of a high-end personal computer (PC) running Microsoft Windows 2000 ProfessionalTM or Windows XP ProfessionalTM. The Master Station software consists of two components:

- Synchronization program The synchronization component of the Master Station software is intended to run automatically without user intervention. Its purpose is to connect to the user's Field Unit(s) automatically on a configurable time interval. Each time it does to, it retrieves any new information that the Field Units have collected. For example, the user may configure his Master Station to run the synchronization program at midnight every night so that any information collected on a given day will be available on the Master Station the next morning when he comes to his office. Because the synchronization program is not normally intended for human interaction, its interface is a simple console that generates text progress messages as it retrieves information from the Field Units. The purpose of the automatic synchronization program is to relieve the user of the burden of manually connecting to Field Units, determining what information has not been collected yet, and going through the tedious process of specifying and retrieving this new information.
- 2. Graphical User Interface (GUI) The graphical user interface, or GUI, component of the Master Station software provides the user with tools to view and analyze both the

Work to Date

statistical information and the data captures that the Field Units collect and that the Master Station synchronization component retrieves. It also allows the user to communicate with the Field Unit in real-time, in order to view real-time readings and to view and modify settings on the Field Unit.

The GUI organizes information hierarchically by utility, substation, bus, and feeder. At the feeder level, it then provides means to filter data captures by date and by cause (e.g., capacitor bank, arcing, overcurrent, etc.). The GUI simplifies visual analysis of data captures by allowing the user to simply double-click on a data capture's icon to open it and display it in the GUI's chart window.

Figure 3-2 illustrates the look and feel of the GUI when viewing one parameter in a data capture. The icon next to the data capture entry in the tree indicates that this particular data capture has been classified as an overcurrent event. Figure 3-3 illustrates the viewing of statistical data, providing a high-level view of the performance of a particular feeder over an extended period of time (e.g., a week, a month, etc.). Note that these screen captures were taken from TAMU's Master Station GUI. Therefore, their hierarchical trees show multiple utility companies and substations. By contrast, each participating utility company's Master Station GUI's only have single Utility entries in their GUI trees.



Figure 3-2 Data Capture Viewed in DFA Prototype Master Station GUI

Work to Date





Initial Problems Encountered and Corrected

As is to be expected with any Prototype system, initial field deployment revealed several shortcomings of the hardware. The following items briefly identify several shortcomings that the users and hardware developer identified and that the developer subsequently corrected:

- 1. Improper soldering on feeder module terminals Late in the delivery cycle, it was determined that the solder joints in the first five Field Units that were produced were of questionable integrity. Because this potentially compromised the continuity of users' CT circuits, the developer requested that the affected users remove their Field Units and return them. The developer inspected the solder joints, reinforced them where necessary, and returned them to the users for reinstallation.
- 2. Hard disk power supply After multiple Field Units operated for several months, a problem surfaced related to the power supplies for their internal hard disks. The power supply voltage dipped during certain sequences of events involving repetitive faults on the monitored power system. When this occurred, the hard disk temporarily shut itself OFF, which in turn caused the Field Unit to perform a "hard" reboot. Interestingly, it later was determined that both the UPS (Uninterruptible Power Supply) that the Field Units use and other mainstream UPS models do not switch to battery backup until their input supply voltage falls to approximately 90-95 volts. Any future commercial design must

Work to Date

take this into account for systems that are to run on AC power. This problem is closely related to item 3 described below and therefore its resolution is described as part of the resolution for that item.

3. Slow boot problem – The Master CPU on the Field Units runs Microsoft Corporation's Windows 2000 Professional[™] operating system. After several months of operation, many of the Field Units began to develop a slow-boot problem. The normal time for the operating system to load and begin running was three to four minutes. A system affected by the slow-boot problem takes 15 to 20 minutes. This problem generally developed after one or more "hard" reboots, such as those caused by the hard disk power supply problem described above. The problem was complicated by the fact that each Field Unit has a watchdog circuit on its Master CPU board. Because these systems must run unattended for extended periods of time, the intent of the watchdog is to reboot the Field Unit if it ever "hangs." The Field Unit's custom software regularly refreshes this watchdog during normal operation, to keep the watchdog from resetting the unit. The time that the watchdog waited between refresh cycles was 15 minutes, with the initial countdown beginning when power is applied to the unit. This was more than adequate for normal boot times of three to four minutes. Unfortunately, when a unit developed the slow-boot problem, the operating system could not load and begin running the custom software quickly enough to keep the watchdog from performing a reboot.

System developers searched for a solution to this problem. They learned that many individuals had encountered this behavior with Windows 2000 ProfessionalTM, but none, including Microsoft technical support personnel, offered solutions other than the default recommendation of "reinstall the operating system." Even if it were reasonable to ask users to reinstall the operating system on their Field Units, this solution is not feasible in the long term. The average time from initial deployment of a Field Unit until it began experiencing this problem was a few months. Therefore, it seems likely that the operating system would have to be reinstalled several times per year, which clearly is unacceptable.

The solution used to correct items 2 and 3 was three-fold: 1) reinstall the operating system one time, 2) add an auxiliary power supply to minimize hard reboots in the future, and 3) install a modified watchdog timer with a 60-minute timeout instead of a 15-minute timeout. These steps were carried out on some Field Units prior to initial delivery, on others when they were returned in response to the problem with solder joints, and on still others while installed in the field.

Users Group Meetings

Active participation by host utility members is critical to this project. These participants are responsible for the day-to-day operation of the DFA Prototypes at their respective utility companies. They also are responsible for investigating and documenting events that occur on the monitored feeders.

It is beneficial for the utility users to interact with one another over the project's duration. In this way, each learns from the experiences of others. Through trial and error, one utility company

may conceive a useful method for diagnosing certain types of problems. By sharing this experience with other project participants, all members build on and benefit from a common body of knowledge. EPRI hosts Users Group meetings approximately every six months for the duration of the project, in order to facilitate focused interaction of this type.

To date, the project has held five Users Group meetings and the sixth is being scheduled:

- 1. July 10-11, 2002, Texas A&M University, College Station, Texas
- 2. March 26-27, 2003, Alabama Power Company, Birmingham, Alabama
- 3. October 15-16, 2003, BC Hydro, Vancouver, British Columbia, Canada
- 4. March 24-25, 2004, City Public Service, San Antonio, Texas
- 5. September 9-10, 2004, Exelon, Chicago, Illinois
- 6. Next meeting: scheduled for Spring 2004, Tennessee Valley Authority

These Users Group meetings have provided a forum for the various participants to become familiar with each other's installations and experiences and have been quite valuable. Most participants have been able to attend the meetings in person, although budget constraints of some of the participating utilities have prevented their participation in-person. In these instances, some of the participants have participated by telephone conferences and WebExTM conferences. Clearly this type of participation is inferior to in-person participation, but it does have value if the only alternative is non-participation.

The first meeting in College Station, Texas was an introductory meeting. TAMU personnel introduced the participants to the project's goals and gave them an opportunity to see an installed DFA Prototype first-hand. They also introduced the users to the look-and-feel of the GUI.

By the time the second meeting was held in Birmingham, Alabama, TAMU had shipped Field Units to all of the initial utility participants and approximately half of the these units had been installed. TAMU provided additional training in the use of the GUI at this meeting.

At the third meeting, in Vancouver, most Field Units had been installed for a significant period of time, although there still were three Field Units (two eight-feeder units and one four-feeder unit, for a total of 20 feeders) that owner utilities had not installed for various reasons internal to their companies. In addition, there were two Field Units (six feeders) that were not installed because the owner utility had joined the project only recently and had not had the opportunity to make arrangements for installation. Users discussed various events that had occurred on their systems. TAMU provided additional training on approaches for recognizing signatures of certain power system events.

The fourth meeting, in San Antonio, Texas, focused on additional faults and other experiences that the users had encountered. The users expressed their feeling of being overwhelmed by the

Work to Date

sheer volume of data that their units were recording on a daily basis, most of which were captures of normal system events like capacitors switching and motors starting. Based upon this feedback, TAMU committed to begin propagating to the users results of their automatic classification algorithms. TAMU's original plan had been to delay doing this until the classifications were more refined and more robust. However, proceeding with their release would provide users a tool that would allow them to analyze their captures much more efficiently. Therefore, it was agreed that TAMU would begin to run classification algorithms automatically and propagate these automatic classifications to the users' Master Stations. TAMU would begin automatic processing with the capacitor-related algorithm first, and then add an additional algorithm approximately every 1-1/2 months.

When the fifth meeting was held, in Chicago, Illinois, TAMU had released two algorithms and the system of automatic classification and propagation of results was going well. The release of these first two algorithms meant that approximately half of the captures on a user's Master Station were being classified automatically. There also was some discussion at that meeting concerning the appropriate timing and method for commercializing the DFA technology.

Automatic Classification Algorithms and System

Many of the events of interest in this project have small magnitudes, often of the same order of magnitude as many loads. Because the magnitudes are similar, there are a very significant number of events each day on the average feeder. To analyze each one manually would be a daunting task, even if one had sufficient time to do so. The reality is that no one would have the time necessary to do this. Therefore, one of the key tasks of this project is the development of methods for automating the recognition of captured data waveforms. Ultimately, the goal is to automate this to the point that the system would not only recognize what particular captures represent, but also be able to decide which ones are "important" and should be brought to the attention of the user.

TAMU has developed and tested algorithms for automatically classifying a large majority of the data captures. So that individual utility users get the benefit of this automatic classification, they also have put in place a system for automatically sending the results of these classifications to the individual utility users. This system also retrieves any classification results that the users have entered into their Master Stations manually.

Conceptually, the automatic classification system consists of two parts: the classification algorithms themselves, and an automated system that retrieves captures, classifies them, and propagates the results to users. The following subsections discuss each of these parts in more detail.

Classification Algorithms

As soon as Prototypes began to be put in place in substations, they started generating a significant volume of data. Normal system operations, such as capacitors and motors switching, are common events, and, as expected, began generating a steady stream of data captures from

day one. Because no fault conditions were staged or otherwise created, faults and fault precursors do not occur with great frequency, so acquisition of significant numbers of these captures did not occur nearly as quickly.

When TAMU began to design algorithms for automating the recognition of data captures, they took into account these realities, and focused on developing the algorithms for normal system events first. There were two complementary reasons for doing this:

- The relatively large volume of data captures representing normal system events soon became sufficiently large to support development and testing of robust algorithms, much more quickly than the volume of data captures representing abnormal events did so. As an example, most circuits have at least one capacitor bank that switches at least twice per day (once turning ON and once turning OFF). This alone generates 30 days/month x 60 circuits x 2 operations/day/circuit = 3,600 captures per month. By contrast, load tap changers (LTCs) seldom fail, so the probability of experiencing even a single LTC failure in a given month (or year, for that matter) is very small.
- 2. One of the responsibilities taken on by each user was the classification of each data capture, by analyzing the waveforms and entering the appropriate classification information into the database via the Master Station software. After gaining experience and proficiency in doing this, they could accomplish the classification of individual normal system captures very quickly. However, the cumulative amount of time to classify a large number of classifications was very significant, particularly for those users who have a large number of monitored circuits or who have particularly active circuits in terms of the frequency of normal events.

When the TAMU research team first designs and implements a classification algorithm, it is normal to expect there to be some number of incorrect results. Therefore, whenever a new classification algorithm was ready for initial testing, they began allowing it to classify new data captures as they occurred and were retrieved, but they did not immediately allow the results of these classifications to be propagated to the users. Instead, over a period of time, they compared the results of the automatic classifications with those manually entered both by utility users and by TAMU personnel. They used this feedback mechanism to assess strengths and weaknesses of the algorithm, and then used this assessment to refine the algorithm. The algorithm's accuracy generally became high enough within the first few weeks of this kind of testing to allow the results to be propagated to the users. Naturally, the accuracy and robustness of a given algorithm was proportional to the volume of data available for testing it. Therefore, an algorithm such as the capacitor-recognition algorithm achieved a very high level of accuracy very quickly in comparison to, say, an algorithm for recognizing the failure of a line switch.

As of the writing of this report, TAMU had "released" several algorithms, meaning that they were not only classifying data captures as they became available, but that they also were propagating these results to the users automatically. As of the writing of this report in early December 2004, TAMU had released automatic classification algorithms for recognizing capacitor operations (both normal operations and numerous types of abnormal operations), motor starts, arcing, and overcurrent faults. The use of the algorithm for classifying overcurrent faults

Work to Date

does two things. Obviously, it provides the basic function of classifying appropriate captures as being overcurrent faults. Less obviously, it provides the basis for recognizing that a fault is happening repetitively. Over the course of this project, various utilities have experienced multiple instances of this, one example of which is described in Chapter 4, Field Results.

Automated Classification System

Currently, the central processing location is a server at TAMU headquarters in College Station, Texas. The automated classification system must perform several tasks in order to make the system work together, without human intervention. To do this, it coordinates the following activities:

- 1. Retrieve new data captures from Field Units.
- 2. Retrieve user classifications from Field Units.
- 3. Run algorithms to classify each new data capture.
- 4. Resolve discrepancies between manual user classifications, manual TAMU classifications, and automatic TAMU classifications, and set icons appropriately so that the filtering functions of the Graphical User Interface (GUI) work properly.
- 5. Propagate appropriate classifications and icons to the Field Units, and ultimately to the utility users' Master Stations.

The TAMU team designed appropriate automation on their Master Station to accomplish these tasks. As of this writing, this system had been in place and working well for several months.

4 FIELD RESULTS

At the time of this writing, many of the DFA Prototypes had been installed for between one and two years. One of the participating utility companies was a relative latecomer to the project and therefore had just installed one of its units, but all of the other units had at least six months of continuous operation, with the majority of them having considerably more.

The following sections provide a summary of collected data and a detailed series of case studies that give a sense of the types of incidents being recorded by the Field Units. The proof-of-concept project that provided the basis for the current project involved a small number of single-feeder pre-prototype units, some of which continued to collect information valuable to this project. Information from these units is included in these case studies.

Some of the cases cited below have been resolved in terms of finding and making repairs to the underlying problems on the power distribution system, but others have not. By their very nature, incipient events often take a considerable amount of time to develop. Therefore, there sometimes is a considerable period of time between initial recognition that something is wrong and final determination of what the problem is.

As often happens in research projects that involve collection of field data, some of the collected information represents events that are somewhat tangential to the initial focus of the project, but that present additional opportunities to use DFA-like monitoring to enhance reliability and operational efficiency.

Data Collection Philosophy

One of the primary goals of this project is to characterize early warning signs that might accompany failing equipment. Because no one ever collected this type of information before, there was no way to know beforehand what form these precursor signals might take. Furthermore, the number of equipment failures per year on a given feeder is small. Another factor that sometimes is frustrating is that neither the research team nor the participating utility companies can cause these failures to occur: they happen when they happen and you'd better catch them when they do!

All of these factors drove the research team's data collection philosophy: record data captures on very small parametric changes, so as not to miss the rare event of true interest when it does come along. Unfortunately, many normal system events affect many of the monitored parameters. Unlike the rare incipient failure, these normal system events occur frequently. For instance, a capacitor bank or a large motor may switch OFF and ON several times per day, each time

causing perturbations in the monitored parameters. Therefore, it is predictable that perfectly normal system events represent the vast majority of data captures obtained in such a project.

Summary of Failures Documented to Date

As mentioned, the data collection methodology is such that, after the Field Unit equipment is installed, neither the researchers nor the utility participants control when failures or incipient failures will occur. As expected from the outset, the process involves a great deal of waiting for power system events of interest to occur naturally. It also involves wading through voluminous normal system data in the quest to identify these infrequent, but most important, failures and incipient failures.

Nevertheless, a wide range of failure modes have been recorded and documented, including one or more instances of each of the following:

- Voltage regulator failure
- Load tap changer (LTC) controller failure
- Tree limbs contacting overhead primary
- Customer service cable (secondary) insulation failure
- Medium (i.e., distribution) voltage cable failure
- URD (underground residential distribution) cable failure
- Customer transformer failure
- Repetitive overcurrent faults
- Cutout and line switch failure
- Numerous modes of capacitor bank failures and operational problems

These are in addition to numerous instances of "normal," unpredictable overcurrents, such as animals on transformers, automobiles hitting poles, etc. These are not included in the list because they are, by definition, unpredictable.

Case Studies

This section details several specific instances of problems on monitored circuits. In some instances, the DFA Field Units have registered multiple instances of a particular variety or mode of failure. In these instances, only one example of that type of failure mode is presented, to avoid redundancy.

Where appropriate, some of the case studies discuss the affected utility's use of information from the DFA, and how they might use that information further if it were integrated into their day-to-day operations.

Malfunctioning Capacitor Bank Controller

From the outset of the project, the presumed failure modes involved mechanical failure of equipment. However, one result of the monitoring activities has been to identify not only these mechanical failures, but also failures in controls that cause the underlying equipment to behave improperly. One of the best examples of this was a malfunctioning capacitor controller.

The utility in question completed installation of its Field Unit, including communications and configuration, in December 2003. TAMU began retrieving capture data from this unit in early January 2004, at which time it was immediately obvious that there was an operational problem developing with one of the capacitor banks on one of the monitored circuits.

On average, one normally would expect a capacitor bank to switch once or twice per day, perhaps with a few more operations on rare occasions. One of the circuits at the utility in question had a capacitor that was switching much more often than this on most days. The number of operations each day was not consistent, but, in general, each day exhibited far more operations than one normally would expect or desire. The following tabulation shows the number of recorded capacitor switching operations, on this one feeder, for two representative one-week periods. Each operation represents either a switch closing or a switch opening.

Date	Capacitor Switching
	Operations
01/01/2004	39
01/02/2004	28
01/03/2004	14
01/04/2004	27
01/05/2004	32
01/06/2004	71
01/07/2004	34
Week Total	245
	(35/day average)
02/01/2004	12
02/02/2004	36
02/03/2004	78
02/04/2004	119
02/05/2004	77
02/06/2004	82
02/07/2004	<u>90</u>
Week Total	494
	(70/day average)

Wow! Something clearly is amiss. For these representative periods, the average number of capacitor switching operations is 35 per day during the first week of January 2004, and 70 per day during the first week of February 2004. The number of operations each day was not monotonically increasing, but there did seem to be a general upward trend. The worst single day was February 18, 2004, when the DFA Prototype recorded 373 capacitor operations on this circuit.

Figure 4-1 illustrates a phase-A capacitor short circuit on February 16, 2004, and represents a catastrophic failure of one of the capacitor "cans" that make up this bank. It is possible that this failure is unrelated to the excessive number of switching operations, but it should be noted that by this time, the DFA Prototype had registered something on the order of 3,000 switching operation for this capacitor in the period of approximately two-and-one-half months since the Field Unit was installed. It seems likely that the resulting repetitive switching transients caused or at least accelerated the failure of this capacitor.



Figure 4-1 Phase Capacitor Short Circuit

Following the February 16, 2004, capacitor failure on phase A, the remaining two phases continued to switch with excessive frequency for the next two weeks. Then, on February 29, 2004, the contacts of the oil switch controlling the phase-B capacitor apparently began to arc. After switching dozens of times in the early morning hours, the bank switched ON at 4:57:37. On this occasion, the contacts of the oil switch controlling the phase-B capacitor made imperfect contact, and, eight seconds later, continuity of the contacts became intermittent.

Figure 4-2 illustrates the RMS current signals the Field Unit measured at the substation when the switch began to fail. Whenever a capacitor switches ON, it causes significant, but generally very time-limited, transient current to flow, as the capacitor quickly charges to the instantaneous line voltage. Intermittent conduction in switch contacts appears electrically like a rapid series of back-to-back opening and closing operations. Therefore, each time the switch begins to conduct, it causes the capacitor to draw the same kind of transient current that occurs when the capacitor switches ON normally. The difference here is that the voltage differential now is between the present line voltage and the voltage left stored on the capacitor the last time the switch stopped conducting, instead of the normal situation in which the voltage of the capacitor is essentially zero when the capacitor switches on. The significance of this is that the transient current will be

less than normal if the polarity of the voltage stored on the capacitor is the same as that of the line voltage at the instant that conduction begins, but greater than normal if the polarity of the stored voltage is opposite that of the line voltage at that instant. In either case, there will be a transient current each time conduction begins. This explains the erratic behavior of the current in Figure 4-2.



Figure 4-2 Continuity Failure in Capacitor Oil Switch Contacts

In addition to drawing transient current each time it switches ON, a capacitor also produces voltage transients, often very significant ones. Those voltage transients also occur each time a faulty switch begins to conduct. During this incident, the faulty switch exhibited countless cycles of losing and then reestablishing contact. As a result, the voltages experienced a nearly continuous series of back-to-back transients. Figure 4-3 shows several cycles of the phase voltage measured at the bus feeding of the affected feeder. The severe transients are obvious and clearly constitute a serious power quality problem. It should be stressed that the figure illustrates the voltage at the bus, meaning that the voltage out on the circuit would be worse than that shown, and also meaning that the other circuit on this bus would be subject to the same repetitive transients.

One might assume that this sort of intermittent making and breaking of the switch contacts would cause the contacts to degrade further, and that, after a short period of time, they would cease conduction altogether. This would leave the circuit with a VAR imbalance, which constitutes an operational inefficiency, but a relatively minor one in comparison with the severe transients that continued, intermittent conduction represent. One might assume this, but, in this case, one would be wrong. In this case, the controller continued to open and close the switch frequently, and whenever the switch was in the closed position, it continued to exhibit these severe, repetitive transients.

Figure 4-4 illustrates the end of the sequence of failures. The intermittent current stopped flowing early on March 4, 2004, nearly four days after the switch contacts began to fail. The data do not make it obvious whether the switch finally stopped conducting, or whether the phase-B capacitor or fuse finally opened in response to the constant abuse over the preceding days. It does not appear that the controller opened the switch at this point, because the other phase (not shown) that still had an intact bank did not register a step in VARs.



Figure 4-3 "Continuous Transients" During Capacitor Switch Failure

TAMU notified the utility about the problem with the overactive controller at the end of January 2004. They also notified the utility of the phase-A capacitor failure on February 16, 2004 (Figure 4-1). Finally, they notified the utility of the switch failure that had gone on for four days, starting February 29, 2004. Because this is a research project and not considered to be part of normal operations, utility personnel took no corrective action until March 9, 2004, at which time they patrolled both affected circuits and took temporary corrective actions. They verified that the failure reports generated from the DFA Prototype data were accurate. They also found that the operations counter on the overactive controller had registered in excess of 4,000 operations since being serviced the previous summer. Interestingly, another utility participant reported that oil switches are supposed to be maintained roughly every 1,500 cycles, and this switch had operated far in excess of this number since the utility last performed service (e.g., contact inspection and replacement, oil replacement) on it.





This is an interesting and important case for multiple reasons. First of all, the repetitive transients generated by the excessive switching operations represent a potential power quality problem. Second, when the switch contacts began to fail, they generated repetitive transients that were experienced by all customers and all equipment on the capacitor's circuit and on the other circuit on the same bus. As a result of these repetitive transients, other capacitors on both circuits failed. In total, the equipment that failed as a direct result of this overactive controller includes two capacitor cans, fuses on two other capacitors, and the oil switch. Obviously, the contacts of the switches on the other two phases also have experienced excessive operations, and, although they were removed from service before they failed, they certainly are due for premature maintenance. Finally, the bus voltages ended up with a significant, sustained imbalance. It turned out that blown capacitors and melted fuses resulted in phase A losing two 400 kVAR capacitors and phase B losing the same amount, but phase C did not lose any of its VAR support. A load tap changer (LTC) on the substation transformer maintains voltage levels at the bus that supplies the two affected circuits. The LTC controller senses one of the phase voltages at the bus and raises all three phase voltages based upon the value of the one that it senses. It is not capable of raising or lowering the voltage of one phase independent of the other two. Therefore, the 800-kVAR imbalance resulted in a sustained voltage imbalance. One of the phase voltages was 4.7 percent higher than the other two, as measured at the substation bus. This voltage imbalance persisted for several days, from the time the failures occurred until the day the utility patrolled the circuits and took corrective action. Finally, it seems likely that the repetitive transients on these two feeders would have done some level of damage to many pieces of connected equipment. Quantifying the amount of damage would be difficult or impossible, but it is safe to assume that there was some level of such damage.

Repetitive Overcurrent Fault

Early on the morning of November 2, 2004, a tree limb came into contact with both a phase conductor and the neutral conductor on an overhead single-phase lateral. At 6:57 that morning, it caused a high-current fault, which was cleared temporarily by a three-phase poletop recloser that tripped and reclosed one time. Figure 4-5 illustrates a second instance, which occurred just over an hour later. In this instance, the high current caused the three-phase poletop recloser to trip, at approximately 2.2 seconds into the time period illustrated in the figure. In this second instance, after the recloser closed back in, the tree limb continued to contact the line and draw significant current. The erratic current starting approximately 4.3 seconds into the time period shown in the figure caused the recloser to trip again at approximately 6.9 seconds. When it reclosed again, the tree limb apparently had lost contact and stopped arcing, at least for the time being.



Figure 4-5 Second in Series of Overcurrent Faults from Tree Limb Contact



Figure 4-6 Final in Series of Overcurrent Faults from Tree Limb Contact

The tree limb did not burn permanently clear of the line, however, because it started faulting again 16 hours later, at just past midnight on November 3, 2004. In total, this tree limb caused 17 high-current faults and recloser operations, at the times tabulated below. When the line burned down, it interrupted power to 140 customers for 62 minutes.

Date	Time	Recloser Trips
11/02/2004	06:57:47	1
	07:58:33	2
11/03/2004	00:09:06	1
	00:16:48	1
	00:40:38	1
	00:40:53	1
	01:10:51	1
	01:12:37	1
	01:15:30	1
	03:24:47	1
	04:19:39	1
	04:30:36	1
	05:51:01	1
	06:19:45	3
	Total Trips	17



Figure 4-7 Tree Limb That Caused Repetitive Faults and Burned Line Down

Figure 4-7 shows two clearly identifiable burn areas on the offending limb. One of the burned areas is near a fork in the tree, which apparently hung on the phase conductor and pulled it down to within 18 to 24 inches of the under-built neutral conductor, bringing the neutral into contact with the second burned area.

Some might question whether this constitutes fault anticipation, since there was no anticipatory signal noted prior to the initial overcurrent, trip, and reclose. In the traditional sense of the word, there is room for debate. However, it is undeniable that this is anticipatory in the sense that some of the later instances of faults and trips, and likely the sustained outage as well, could have been prevented. As a follow-up, TAMU asked the engineer responsible for the DFA project at this utility how he would have used this information, had he had real-time access to it. His responses are summarized as follows:

- 1. We definitely would have dispatched a crew to investigate around 1:00 AM, given the information that the DFA provided.
- 2. I believe it likely that the crew would have found the fault in time to prevent the line from burning down and causing the ultimate outage.

3. I believe that we could have found the fault by using the following information, based upon data from the DFA Prototype: the phase of the fault, the magnitude of the fault, and the recloser's operating characteristics.

Long-Term, Repetitive Arcing

Several of the monitored circuits register sporadic arcing signatures from time to time. A circuit may produce a singular arcing episode that lasts from a cycle or so to a few seconds. In some of these instances, the arcing never appears to recur or cause any problem. In other instances, a particular circuit may begin to produce repeated episodes of arcing. The pattern may repeat itself numerous times and eventually may escalate to something more serious.

An interesting series of arcing episodes is being investigated at the time of writing of this report. All arcing captures involved phase B and ground and each produced a few tens of amperes for durations ranging from one cycle to approximately one second. There was no clear pattern to the times of occurrence of these episodes of arcing. Figure 4-8 shows one of the early episodes and Figure 4-9 shows one of the last.



Figure 4-8 Early Episode of Repetitive, Intermittent Arcing



Figure 4-9 Episode of Repetitive, Intermittent Arcing, Almost a Month Later

In total, as of the writing of this report, the DFA Prototype had recorded one or more episodes of arcing on this circuit on each of the following nine dates:

Date	Episodes Recorded
11/11/2004	3
11/15/2004	6
11/24/2004	4
11/28/2004	4
12/02/2004	4
12/03/2004	4
12/04/2004	1
12/06/2004	5
12/07/2004	<u>1</u>
Total Episodes Recorded	32

When TAMU first contacted the utility about this in mid November 2004, they were unaware of any trouble that they had experienced related to these early instances. As more episodes continued to occur over time, they patrolled the feeder to look for vegetation-related problems. They found two places where vegetation was a potential problem, one of them involving a tree limb in the line and the other involving vines in the line. They cleared these problems during the patrol period between December 3 and 7. As is apparent in the tabulation above, removal of these vegetation intrusions did not remedy the problem, because the arcing recurred on the morning of December 7, 2004. As an aside, the patrol crew also found and fixed a wire that had come loose from its porcelain insulator and was lying on its crossarm. However, this was not considered as a potential candidate for the repetitive arcing problem, because it was on a different phase.

The utility has a location in which they serve four customers using direct-burial, underground secondary service cable. A pole-mounted 50-kVA service transformer feeds this underground cable. The service cable comes from the transformer secondary bushings, down the pole, and into a buried connection box. This box contains a "crab" that connects the four customers' service cables to the cable from the transformer. Individual service cables then go from the box to each customer's service entrance.

The utility knew that there was some kind of problem at this location, because the service transformer had tripped several times in recent weeks. Each time the transformer tripped, troublemen looked for a cause but found none. When they reset the transformer, it restored service, so they left the location. They believed that there was some sort of problem, but they were unable to determine what it was.

The transformer tripped again early on December 7, 2004, generating an trouble report at a time that corresponded within minutes to the time of the arcing episode of Figure 4-9. The utility engineer responsible for the DFA project was informed of this outage, and the coincidence of times became obvious. After making this logical connection between the outage and the DFA records, the utility looked back at the other recent times at which this service transformer had tripped and determined that this transformer had tripped three times between December 1 and 7, 2004, and each of the times was coincident to within minutes of recorded arcing episodes reported by the DFA Prototype.

The utility believes it likely that the problem is in the one of the secondary service cables, as opposed to being in the connections in the buried box. They believe this because the local gas company recently did some excavation work near three of the service cables, and the utility believes that this work may have disturbed, and possibly even damaged, one or more of the customer service cables. At the writing of this report, neither the buried box nor the underground cables was available for inspection. The utility company was in the process of digging up the box and, perhaps, the cables, for inspection. This potentially will provide additional valuable information, because it should provide even more detail about the exact nature of the failure.

After the source of the problem was verified, Texas A&M asked the engineer responsible for the DFA project at this utility how that utility would use the information from the DFA in a future situation like this, if the information were available and integrated into their day-to-day operations. His responses were as follows:

- 1. "We are already using the information [from the DFA] to investigate the cause of the repetitive arcing. Phase identification is beneficial in that it reduces the number of phases that must be tracked when investigating or 'riding' a feeder to locate the problem. This is very beneficial on a radial overhead system. Tracking the times allows us to verify the report with other sources thereby validating the occurrence and justifying the need to determine the cause, (i.e., avoid potential outage while increasing customer satisfaction with performance)."
- 2. "Action would have been taken sooner had we been a little more sure of what it was we were looking for. Narrowing down the possibilities of arcing allows us to more closely

focus on what the problem might be. It's much easier to look for the characteristics of tree contact then for a failing insulator or lightning arrestor. The DFA has already identified an area ... [location omitted for confidentiality] that is subject to rapid tree/vine growth and, as a result, the area is now scheduled for more timely and aggressive trimming. The inclusion of the DFA as a means of identifying possible problems and feeder outages has already been incorporated here! What we hope is the DFA will be an expandable problem solver not only able to help eliminate potential problems before they become unmanageable but a tool to help reduce costs by allowing us to put our maintenance money where it will do most good."

Failing Switches

At 11:57 AM on March 16, 2004, one of the Field Units began to register signals similar to the one shown in Figure 4-10. The Prototype registered these signals nearly continuously for the next 42 minutes, as a result of a fused cutout that was beginning to fail.

When a crew arrived at 12:39, the top of the pole was engulfed in flames. The crew manually tripped the substation breaker to isolate the fault. They switched out the affected section of line and restored service to the rest of the circuit. They then repaired the top of the pole, including replacement of the cutout, and restored full service at 2:29 PM.



March 16, 2004 11:57:52

Figure 4-10 Measured Signals Produced by Failure of Cutout

This incident is of special interest, because it illustrates one of the inherent facets of a major research project that involves field measurement. The utility followed up on the incident shortly after it happened, and the relevant data captures were labeled as being associated with a cutout failure that caused a pole top fire. However, it was not until approximately six months later that the TAMU research team identified what they believed to be a characteristic signal unique to failing switches. After identifying this characteristic, they were able to review this and other "old" data captures that previously had not been identified. From this review, they were able to

identify four instances of apparent switch failures. They then requested that the corresponding utilities perform additional reviews of their trouble records at the times of these incidents. In every identified case, they found that either a line fuse or a cutout failure had occurred. The point is that sometimes a "break-through" case is needed in order to identify a characteristic unique to a particular failure mode. When this occurs, the vast library that this project is creating allows review and validation of the finding, even when the relevant data captures originally were not recognized at the time they occurred.

Short-Lived, High-Current Events

It is tempting to assume that, in general, a high-current event (e.g., one with peak currents in excess of 1,000 amperes) will not extinguish itself, but rather that it will persist until some outside force interrupts it. This does not mean that such an event necessarily would result in a sustained outage, but that it would require at least momentary interruption, after which successful reclosing might be possible. The authors of this report certainly assumed that this generally would be the case. However, as with many assumptions, this one turned out to be incorrect.

In May 2004, the research team noted the high-current event shown in Figure 4-11. The peak current during this fault, as measured at the substation, was approximately 2,400 amperes, and its duration was between 1/4 and 1/2 of one cycle. Faults of this general magnitude and duration are fairly common, often having as their cause short circuits that clear small fuses on customer service transformers in response to, for example, animal contacts. However, in this case, the utility could find no record of any outage associated with this fault. In addition, examination of the captured data waveforms offered no evidence of a momentary interruption by a recloser. In other words, it appeared that the fault truly self extinguished, even though this was contrary to the prior assumption that high-current events generally would not self extinguish.

Figure 4-12 shows the current produced by a similar short-lived fault that occurred on the same circuit. The fault was very similar to the one that occurred days earlier, in that both occurred on the same phase, both had approximately the same peak magnitude, and both had approximately the same duration. In addition, neither caused an outage or even a momentary interruption.

After noting these occurrences, the research team started looking more closely for other instances of this type of behavior, both at this utility and at others. They were surprised to find that many of the utilities had significant numbers of these high-current, but short-lived, faults, for which there were no records of outages, and for which the data contained no evidence of any momentary interruption or reclosing.

The research team currently has several working theories about possible causes of these shortlived, high-current, self-extinguishing faults. To date, these theories have been neither proved nor disproved. TAMU continues to monitor these types of events with special interest.



Figure 4-11 First Instance of High-current, Short-duration Fault



Figure 4-12 Second Instance of High-current, Short-duration Fault

Deceptive Capacitor Fault

During the course of this project, DFA Prototypes have registered a significant number of capacitor failures of various varieties. Capacitor fuses can blow for several reasons. A high-current failure generally means either that the capacitor has developed an internal short circuit or that there is an external problem, perhaps a cracked bushing. Another mode of fuse operation is one in which the fuse simply opens, without any associated high-current event. This can mean that a long-term condition such as excessive harmonics is causing excess heating of the fuse, or it can mean that the capacitor has developed a partial internal failure that is causing it to draw more steady-state current than normal.

Figure 4-13 illustrates a high-current capacitor fuse operation on August 25, 2004. TAMU notified the affected utility of this apparent capacitor failure shortly thereafter. The utility confirmed that the fuse was open and dispatched a crew for repairs on September 1, 2004. The crew tested the capacitor using standard test procedures. The capacitor tested "good," so the crew simply replaced the fuse, but replaced the existing 50-amp fuse with a 60-amp fuse. When they energized the capacitor, the fuse held and the crew left it in operation.

This seemed odd, because this represented the first instance in which a high-current operation of a capacitor fuse was not related to a permanent problem, such as a failed capacitor. Two days later, the plot thickened, when the same capacitor experienced another fuse-blowing short circuit, as shown in Figure 4-14. The only significant difference between the first and second instances was that the second caused significantly higher fault current, possibly because the larger fuse size allowed it to conduct longer before clearing.

The utility confirmed the second fuse blowing shortly after it occurred, but they did not effect a repair until early November 2004. Surprisingly, when the crew tested the capacitor, it again tested "good," and they placed it back in operation on November 10, 2004, this time with a 75-amp fuse. As of December 6, 2004, the capacitor remained in service without further failure, but both TAMU and the utility continue to watch this situation with interest.

Of note, the second fault removed the capacitor from service a mere two days after the first fuse was replaced. Had this replacement been done as part of normal annual maintenance, the subsequent failure two days later likely would not have been discovered until crews performed maintenance again a year later, meaning that the capacitor would have been out of service all year, despite the fact that normal annual maintenance was performed.



Figure 4-13 First Instance of Capacitor Fault



Figure 4-14 Second Instance of Fault of Same Capacitor

Arcing on Secondary Customer Service Cable

One of the DFA Prototype Field Units measured a series of events on one instrumented feeder over a period of approximately two weeks. The recorded waveforms gave very clear indications of low-current arcing. In addition, the similarities between the several recordings made it almost certain that all of the incidents had the same cause. The recorded waveforms also made it apparent that whatever was causing the arcing was not of sufficient magnitude or duration to operate conventional protection, except perhaps a very small fuse. Further evidence of this was the fact that the incidents continued to occur over an extended period of time.

Figure 4-15 shows the phase current waveform contained in the data capture that the Field Unit recorded during one such incident. Figure 4-16 shows the same time period and illustrates the high-frequency current measured, after being processed with a proprietary algorithm to remove background noise. Both figures include the phase voltage waveform as a reference.

These figures illustrate but one measured instance of the event. The Field Unit made similar captures at each of the following times:

2003/05/31	03:00:48
	03:56:28
2003/06/01	11:58:56
	12:08:01
2003/06/04	05:00:43
2003/06/09	22:12:27
2003/06/10	13:33:04
2003/06/14	21:43:59
	21:44:06
	21:44:12
	21:44:22
	21:44:50 (fuse operation)
	22:14:36 (service restoration)

The last two listed events represent the time at which the arc current finally blew a fuse, interrupting 27 customers, and the time at which service was restored.

This event is of significant interest to the project's goals in that it represents the recognition of insulation deteriorating over a period of weeks. This provides the basis for notifying the utility of the problem while it is in its incipient stage.



Figure 4-15 Arcing Measured on Secondary Service Cable



Figure 4-16 High-frequency Current from Arcing on Secondary Service Cable

Sticking Load Tap Changer (LTC) Controller

The substation transformer feeding a circuit monitored by a Pre-prototype Field Unit began to experience repetitive sequences of events in which the bus voltage stepped down multiple times within a short period of time. Such a sequence of events, in itself, may represent normal operation of the transformer's load tap changer (LTC), for instance after a large load or a large switched capacitor bank switches. However, in many of these instances, after the voltage stopped stepping, the bus voltage was well below the nominal range of acceptable voltages for that bus. In addition, in many instances, the voltage made multiple steps back up less than a minute later.

Sequences of multiple voltage steps were recorded at each of the following times:

2003/02/04	02:52:11 - six steps down (7384 initial voltage)
	02:52:55 - four steps up (7426 final voltage)
2003/02/05	05:29:07 - five steps down (7353 initial voltage)
	05:29:50 - five steps up (7419 final voltage)
2003/02/16	07:57:15 - six steps down (7400 initial voltage)
	07:57:58 - five steps up (7466 final voltage)
2003/02/27	22:49:50 - two steps down (7360 initial voltage)
	22:50:32 - one step up (7451 final voltage)
2003/03/04	02:26:50 - six steps down (7404 initial voltage)
	02:27:36 - four steps up (7474 final voltage)
2003/03/06	03:51:02 - six steps down (7352 initial voltage)
	03:51:44 - five steps up (7436 final voltage)
2003/03/07	05:19:01 - five steps down (7377 initial voltage)
	05:19:45 - three steps up (7379 final voltage)
2003/03/07	20:45:40 - three steps down (7401 initial voltage)
	20:46:28 - one step up (7444 final voltage)
2003/03/10	22:37:21 - three steps down (7375 initial voltage)
2003/03/11	02:00:00 - six steps down (7361 initial voltage)
	02:00:44 - four steps up (7427 final voltage)
2003/03/19	23:55:33 - three steps down (7301 initial voltage)
2003/03/22	00:41:35 - five steps down (7361 initial voltage)
	00:42:20 - four steps up (7455 final voltage)
2003/03/22	09:27:01 - four steps down (7453 initial voltage)
	09:27:47 - three steps up (7450 final voltage)
2003/03/23	16:13:29 - two steps down (7380 initial voltage)
2003/03/27	05:56:00 - nine steps down (7390 initial voltage; 6982 final)
	06:36:05 - ten steps up (7021 initial voltage; 7435 final)
2003/03/29	02:19:59 - three steps down (7341 initial voltage)
	02:20:44 - two steps up (7423 final voltage)

This list shows that this sequence of multiple steps down followed by multiple steps back up was measured intermittently over a period of approximately two months. Most of the sequences occurred in pairs, i.e., a sequence of steps down followed less than a minute later by a sequence of steps back up.

Figure 4-17 illustrates a special case in this sequence of events. At 5:56 AM on March 27, 2003, the bus voltage stepped down, similar to the way it had done many times before, except that the number of steps was greater than usual. In this case, the bus voltage stepped down nine times over a period of sixteen seconds. However, unlike the previous cases of downward steps, in this instance there was not a sequence of upward steps shortly thereafter. The next measured voltage steps up occurred 39 minutes later, at 6:35 AM, when the voltage took ten upward steps.





Around 6:00 AM on this date, a customer reported a low-voltage condition. The utility dispatched a trouble man to investigate. He found that the circuit breaker on the LTC controller motor was tripped. Not knowing why the breaker had tripped, he reset it and the LTC stepped back up to a normal value at 6:35 AM.

This customer call provided the utility with an indication of a problem, but field personnel did not have the information cited above, and, thus, had limited means to diagnose the problem's underlying cause. When the cited information was made available to them, personnel investigated the most obvious cause of the problem: that the LTC controller was sticking intermittently. They found that the LTC controller's manufacturer had issued a service bulletin regarding the need to install an auxiliary component in the controller to keep its mechanism from sticking and causing this type of improper operation. Following the service bulletin's recommendation cured the problem.

This case illustrated the ability of DFA technology to provide information critical to the efficient diagnosis of malfunctioning equipment on distribution circuits.

Inoperative Capacitor Bank Switch

For decades, most utility companies have placed fixed and switched capacitor banks on their feeders in order to provide voltage and VAR support. Switched banks have control algorithms that use a variety of parameters, including time-of-day, temperature, voltage, etc, to determine when the capacitor should be switched OFF and ON. Although there are some exceptions, most switched capacitor banks are operated on a balanced, three-phase basis. In other words, most

banks are configured with three equally sized units and their controllers switch the three phases ON and OFF at the same time.

There are several ways in which capacitor banks fail to operate properly. One of the most common failure modes is the blowing of a single-phase fuse, often because the capacitor has developed a short circuit internally. Another is the failure of a switch to open or close properly. The result of either of these failure modes is the same: the loss of voltage and VAR support on the affected phase. In addition, single- or two-phase operations cause unbalanced voltages, elevated neutral currents, and other operational problems.

As part of their normal operation, DFA Prototype Field Units regularly record data when capacitors switch OFF and ON. Many of the monitored feeders generate multiple capacitor-related data captures each day. The following series of figures illustrates two such captures, the first representing a capacitor switching ON normally and the second representing the same bank switching ON abnormally.

Figure 4-18 shows the VARs recorded when a timed capacitor bank switched ON just before 5:00 AM on June 28, 2003. The VARs on each phase decrease precipitously by approximately 400 kVARs per phase. The three phases do not switch at exactly the same instant, but they all switch within a second or so of one another, so the operation is considered normal.

Figure 4-19 illustrates the same capacitor bank switching on three days later. As in the previous instance, VAR flows drop by approximately 400 kVARs per phase. However, in this case, the VARs on only two phases change, while the third does not. Project personnel noted this change in behavior beginning on June 29, 2003, and reported it to utility field personnel on June 30, 2003. A line crew visited a 1.2 MVAR capacitor bank on the affected feeder and found that one of the oil switches responsible for switching the capacitor bank ON and OFF was not functioning. They replaced the switch within a week of its first failure to operate, restoring proper operation.

Figure 4-20 illustrates the RMS currents measured coincident with the VAR flows shown in Figure 4-19. Of note, the neutral current increased from a level of 18 amps before the capacitor bank switched on to a level of 38 amps after the unbalanced switching occurred. Also of note, phase A did not enjoy the reduction in current that the other two phases experienced.

This case is interesting because it demonstrates the ability to recognize common capacitor bank problems very soon after they develop. Without such monitoring, the capacitor bank would switch into this unbalanced condition daily until the utility company's next inspection cycle. Utility participants in this project indicate that most utility companies perform routine inspections and testing of capacitor banks at specific intervals, ranging from once per year to four times per year, with the most common practice being to inspect and test each bank once per year. Therefore, statistically, an unbalanced capacitor condition such as this would remain undetected, on average, for six months. For this period of time, the utility company and its customers would not enjoy the voltage and VAR support for which the utility originally purchased and installed the capacitor bank. Further, for this period of time, the unbalanced condition would cause some degree of voltage imbalance and increased neutral current.

Between 10 and 15 percent of the feeders monitored by Field Units exhibited at least one capacitor-related problem during the first few months of monitoring. Some of the problems existed when Field Units first started monitoring particular feeders, indicating that the problems had existed for unknown periods of time prior to the installations.



Figure 4-18 VARs Recorded When Balanced Capacitor Bank Switched ON



Figure 4-19 VARs Recorded When Unbalanced Capacitor Bank Switched ON



Figure 4-20 RMS Amps Recorded When Unbalanced Capacitor Bank Switched ON

Temporarily Stuck Capacitor Bank Switch

This case study is interesting because it represents an intermittent condition that affects operational efficiency. Figure 4-21 illustrates the three-phase VAR flows on one of the monitored feeders for the period March 30 – April 17, 2003. For the first two days during the illustrated interval, and again for the last two days of the interval, the three-phase VAR flows appear fairly well balanced and they tend to track one another. However, for the period April 2-14, 2003, there is an obvious deviation from this well-behaved pattern. Specifically, on April 2, 2003, a capacitor apparently switched ON. Later in the day, it switched back OFF. However, when it did so, one phase remained energized, creating a VAR imbalance of approximately 400 kVAR between phases. Because the bank's controller uses temperature as a control parameter, and because temperatures in the area were moderate in April, the bank did not switch ON again until April 14, 2003. When it did, it restored the relative balance between the three-phase VAR flows. Later that day, the capacitor bank switched OFF again, and this time all three phases properly switched OFF.

This case is particularly interesting because of the intermittent nature of the problem. In the previous case, once the problem developed, the failed switch consistently failed to operate. Therefore, although the next inspection and testing cycle might be months away, once it arrived it almost certainly would identify the problem so that it could be repaired. The present case, however, presents a less certain outcome. If a switch only operates incorrectly some of the time, what will it do when a crew performing routine testing operates it? It might fail, in which case the crew can fix it or replace it. Alternatively, it might operate correctly during testing, only to fail again the next time it is called upon to perform its normal function. In such a scenario, its operation likely will be intermittent until the next inspection and testing cycle, often a year later.



Figure 4-21 VAR Flows Caused by Intermittently Sticking Switch

Intermittent Capacitor Switch Restrike

Many of the Field Units capture data daily when capacitor banks switch OFF and ON. When a capacitor is switched ON, a relatively large transient affects the voltage and the current, because this switching operation attempts to change the previously unenergized capacitor voltage to line voltage instantaneously. When a switch is opened to disconnect the capacitor from the line, a similar phenomenon generally does not occur. The switch opening operation breaks the connection between the line and the capacitor, leaving a significant charge on the capacitor, but because the operation does not attempt to change the capacitor's voltage instantaneously, there generally is no significant transient.

One of the monitored feeders began producing significant transients intermittently when one of its capacitor banks switched OFF. Figure 4-22 illustrates the current signals produced during one such instance. The illustrate waveforms have been processed through TAMU's proprietary algorithm for removing the steady-state components.

In the figure, the 300+ ampere spike is obvious. What is less obvious is that a fraction of a cycle before this, there is a steady-state change of approximately 30 amperes. This increase begins at the time at which the capacitor switch opened. The 30-ampere waveform represents the amount of capacitive current that is lost to the system when the capacitor is removed. Figure 4-23 provides a more detailed view of the moment at which the transient occurs and shows the phase voltage waveform as well.



Figure 4-22 Restrike Current When Capacitor Bank Switched OFF



Figure 4-23 Restrike Current and Voltage When Capacitor Bank Switched OFF

Analysis of the facts and waveforms surrounding these captures indicates that the capacitor switch is restriking intermittently when called upon to remove the capacitor from the line. In Figure 4-23, note that the voltage is near its positive peak at the moment the switch opens, as indicated by the fact that this is the time at which the current signal becomes non-zero. (Remember, the current signal represents the amount of capacitive current lost to the power system, so it becomes non-zero when the capacitor switches OFF.) The capacitor is left with a retained voltage equal to the line voltage at the instant that the switch opened, which in this case is nearly equal to the line voltage peak of over 28,000 volts. The capacitor can retain this voltage

for quite some time, so as the line voltage progresses to zero and then becomes negative, the voltage across the switch contacts becomes large. When the line voltage reaches approximately -19,000 volts, the retained capacitor voltage still is +28,000 volts, producing a voltage differential across the switch contacts of approximately 47,000 volts. At this point in time, less than 10 milliseconds have elapsed since the switch contacts opened, so they have had little time to achieve significant mechanical separation. The high voltage across the relatively short physical gap causes the switch to restrike.

Over a four-week period beginning October 2, 2003, this capacitor switched OFF 37 times. Of these, 18 exhibited the restrike phenomenon, always on the same phase. The remaining 19 did not exhibit signs of restrike. Each time the restrike occurred, the waveforms were very similar: The capacitor current transient occurred very shortly after the capacitor switched OFF and had a peak magnitude around 300 amperes.

What is the relevance of this? Does an occasional restrike indicate a problem with the switch? Are several restrikes per week excessive? Does the utility need to take action? If so, does it need to be immediate? These all are questions that currently do not have answers and that ultimately may need to be decided by utility companies. As this project continues, however, future results may provide more insight.

5 PRESENT STATUS AND FUTURE WORK

DFA Prototypes have been designed, constructed, and installed in electric power utility distribution substations. At the time of publication of this interim report, they will have collected over 1,000 feeder-months of data from numerous utility companies with diverse operating conditions and philosophies. Over this period of time, they have collected thousands of capture files. The vast majority of these represent normal system events such as capacitors switching, motors starting, etc. This is by design: The characteristics of most events of interest are, by definition, not known before the Field Units have recorded them and personnel have characterized them. Therefore, collection of a large number of normal events and even "non events" is necessary in order to minimize the likelihood of missing the infrequent, but very important, events of real interest.

The library of recorded events now includes a number of abnormal system events and this library continues to grow with the passage of additional monitoring time. Many of these events have been diagnosed and documented by the utility companies, providing the basis for Texas A&M's research team to develop and test algorithms for classifying these types of events.

The library of captured data contains several instances of data captures that obviously represent abnormal system operations, but for which there is no clear explanation to date. Texas A&M recently had a "breakthrough" event recently, in which the ability of one utility to determine the cause of a particular capture then allowed several "old" events to be properly classified, and verified by the affected utilities, because the old data captures shared common characteristics with the more recent one. Texas A&M believes it likely that other "breakthrough" events will occur as time goes on, allowing an ever-increasing number of these initially unrecognized data captures to be pinned down.

Texas A&M personnel have developed and put into practice methods for classifying certain subsets of the captured events. Participating utility companies currently are seeing benefits from this, in that the results of these automatic classifications are being propagated, automatically, to users on a daily basis, reducing their burden when manually classifying the events on their circuits, and drawing their attention to some of the events of particular interest.

During these first two years of field monitoring, it has become evident that the presence of such an advanced data collection platform, installed in such a variety of locations across the continental United States and Canada, provides the basis for other, related areas of work in the future. It is likely that continued monitoring will suggest additional areas of exploration.

Export Control Restrictions

Access to and use of EPRI Intellectual Property is granted with the specific understanding and requirement that responsibility for ensuring full compliance with all applicable U.S. and foreign export laws and regulations is being undertaken by you and your company. This includes an obligation to ensure that any individual receiving access hereunder who is not a U.S. citizen or permanent U.S. resident is permitted access under applicable U.S. and foreign export laws and regulations. In the event you are uncertain whether you or your company may lawfully obtain access to this EPRI Intellectual Property, you acknowledge that it is your obligation to consult with your company's legal counsel to determine whether this access is lawful. Although EPRI may make available on a case by case basis an informal assessment of the applicable U.S. export classification for specific EPRI Intellectual Property, you and your company acknowledge that this assessment is solely for informational purposes and not for reliance purposes. You and your company acknowledge that it is still the obligation of you and your company to make your own assessment of the applicable U.S. export classification and ensure compliance accordingly. You and your company understand and acknowledge your obligations to make a prompt report to EPRI and the appropriate authorities regarding any access to or use of EPRI Intellectual Property hereunder that may be in violation of applicable U.S. or foreign export laws or regulations.

The Electric Power Research Institute (EPRI)

The Electric Power Research Institute (EPRI), with major locations in Palo Alto, CA, and Charlotte, NC, was established in 1973 as an independent, nonprofit center for public interest energy and environmental research. EPRI brings together member organizations, the Institute's scientists and engineers, and other leading experts to work collaboratively on solutions to the challenges of electric power. These solutions span nearly every area of power generation, delivery, and use, including health, safety, and environment. EPRI's members represent over 90% of the electricity generated in the United States. International participation represents over 10% of EPRI's total R&D program.

Together...shaping the future of electricity

Program:

1001688

Distribution Operations

© 2005 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute and EPRI are registered service marks of the Electric Power Research Institute, Inc.

Printed on recycled paper in the United States of America