

SAND REPORT

SAND2005-7301 Unlimited Release
Printed December 2005

Analysis of Operations and Cyber Security Policies for a System of Cooperating Flexible Alternating Current Transmission System (FACTS) Devices

Laurence R. Phillips, Michael Baca, Jason Hills, Jonathan Margulies, Bankim Tejani,
Bryan Richardson, and Laura Weiland

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.doe.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



Analysis of Operations and Cyber Security Policies for a System of Cooperating Flexible Alternating Current Transmission System (FACTS) Devices

Laurence R. Phillips, Michael Baca, and Bryan Richardson
Critical Infrastructure Systems Department

Jason Hills
Information Operations Red Team and Assessments Department

Jonathan Margulies
Network Systems Survivability and Assurance Department

Bankim Tejani
Knowledge Discovery and Extraction Department

Laura Weiland
System Sustainment and Readiness Technologies Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-1368

Abstract

Flexible Alternating Current Transmission Systems (FACTS) devices are installed on electric power transmission lines to stabilize and regulate power flow. Power lines protected by FACTS devices can increase power flow and better respond to contingencies. The University of Missouri Rolla (UMR) is currently working on a multi-year project to examine the potential use of multiple FACTS devices distributed over a large power system region in a cooperative arrangement in which the FACTS devices work together to optimize and stabilize the regional power system. The report describes operational and security challenges that need to be addressed to employ FACTS devices in this way and recommends references, processes, technologies, and policies to address these challenges.

— This page intentionally left blank —

Contents

1	Introduction.....	9
2	Power Transmission System Reference Model and Context.....	11
3	Description of a System of Cooperating FACTS Devices.....	15
3.1	Standalone FACTS Function and Operations.....	15
3.1.1	Voltage-Sourced Converter.....	15
3.1.2	Static Synchronous Shunt Compensator (STATCOM).....	15
3.1.3	Static Synchronous Series Compensator (SSSC).....	15
3.1.4	Unified Power Flow Controller (UPFC).....	15
3.2	Assumptions about Cooperating FACTS Device Operation.....	16
3.3	Hardware/Software Integrated System.....	16
3.3.1	Algorithms of Interest.....	16
3.3.2	Hardware Elements.....	17
3.3.3	Software Elements.....	18
3.4	Cooperating FACTS Function and Operations.....	19
3.5	Computation of long-term setpoints.....	20
4	FACTS Information Flows & Repositories.....	23
4.1	Switch-Level Control Interactions.....	24
4.2	Dynamic Control Interactions.....	25
4.3	Long-Term Control Interactions.....	25
4.4	Software Code Elements.....	26
4.5	SCADA Interactions.....	26
4.6	Special Challenges in a CFD Environment.....	27
4.6.1	Recognizing and Responding to Contingencies.....	27
4.6.2	Data Completeness.....	28
5	Using Agents to manage FACTS devices.....	31
5.1	Consideration of Agent technology for Electric Power Management.....	31
5.2	Environment.....	32
5.3	Roles.....	32
5.4	Interactions.....	33
5.5	Organizations.....	33
5.6	Liveness and Safety Conditions.....	33
5.7	Conclusions about the CFD Agent System.....	36
6	Consequences of Failure in an Electric Power System.....	37
6.1	Consequence Metrics.....	37
6.2	Public Consequences.....	37
6.3	FACTS-Specific Consequences.....	38
6.4	Power System Consequences.....	39
7	Contingency Analysis.....	41
7.1	Contingency Analysis Types.....	41
7.2	Anatomy of Contingencies and Failures.....	42
7.2.1	Anatomy of a major regional blackout:.....	42
7.3	Natural Vectors for Power System Contingencies.....	43
7.3.1	Normal Load Changes.....	43
7.3.2	Generator Failures.....	44
7.3.3	Transmission System Failures.....	44
7.3.4	Control System Failures.....	44

7.4	Adversarial Vectors for Power System Contingencies	44
7.4.1	Denial of Cooperative Operation	44
7.4.2	Spoofing.....	45
7.4.3	Gaining Control over a CFD or Control Device	45
7.4.4	Desynchronization and Time-Based Attacks	45
7.4.5	Data Injection.....	45
7.4.6	Malware Injection.....	45
7.4.7	Social Engineering.....	46
8	FACTS Security Analysis	47
8.1	Assumptions for this Analysis.....	47
8.2	Security-Related Definitions	47
8.2.1	Security Element Definitions	47
8.2.2	Network type definitions	51
8.3	A System of Cooperating FACTS Devices as a Target of Evaluation	51
8.4	FACTS Information Assets	52
8.5	Inherited SCADA Security.....	53
8.5.1	Physical Layer Communications.....	53
8.5.2	Control Architecture.....	55
8.6	Vulnerabilities.....	56
8.6.1	Operating System and Application Vulnerabilities.....	56
8.6.2	Timing Errors/Induced Delays.....	56
8.6.3	Code Revision/Replacement.....	57
8.6.4	Trust Manipulation.....	57
8.7	Good Security Practices	57
8.7.1	Firewalls.....	57
8.7.2	Access Control Lists (ACLs)	58
8.7.3	Intrusion Detection	58
8.7.4	Authentication.....	58
8.7.5	Logging.....	58
8.7.6	Tamper Prevention	58
8.7.7	Encryption.....	58
8.7.8	Integrity Checking.....	58
8.7.9	Configuration Management	59
8.7.10	Non-Repudiation	59
8.7.11	Redundancy.....	59
8.8	Agent-Based Security.....	59
8.8.1	Control Architectures for Agent-Based Operation.....	59
8.8.2	Security-Oriented System Behavior	60
8.8.3	Multiparty Strategies	62
8.8.4	Agent-Based Mitigation Strategies.....	63
8.9	Network Situational Awareness and Visualization.....	63
8.10	Security Documentation	65
8.10.1	Security Policy	65
8.10.2	Other Important Security Documentation.....	66
9	Conclusions	69
10	References	71
11	Definitions	75

Figures

Figure 1. Information flows in and around an individual FACTS device	19
Figure 2. Notional FACTS layout and its relationship to the power system	23
Figure 3. Switch-level architecture of the UPFC FACTS device.....	24
Figure 4. Maximum Flow Algorithm Information Flows.....	25
Figure 5. Illustration of Power System Transitions through various Contingencies.....	43
Figure 6. The Client-Server Architecture has a central information source, i.e., the server...	56
Figure 7. The Peer-to-Peer Architecture has no central information source; devices can get data only from one another.	59
Figure 8. Location of Vulnerabilities in a CFD Environment.....	62
Figure 9. AnVEBIDAS screen showing the state of the IEEE 118-bus system.....	64
Figure 10. Detail of the AnVEBIDAS view of the IEEE 118-bus system	64
Figure 11. Sandia National Laboratories Security Policy Framework	66
Figure 12. Relationships within the Security Administration Hierarchy.....	67

Tables

Table 1. Switch-Level Control Interactions I/O	25
Table 2. Long-Term Control Interactions I/O	26
Table 3. SCADA Interactions I/O.....	27
Table 4. Categories of attacks on CFD devices.....	49
Table 5. Human Threats.....	50
Table 6. Types of information exchange in a CFD Environment	52
Table 7. Enumeration of Reporting Possibilities.....	60

— This page intentionally left blank —

Analysis of Operations and Cyber Security Policies for a System of Cooperating Flexible Alternating Current Transmission System (FACTS) Devices

Laurence R. Phillips, Michael Baca, Jason Hills, Jonathan Margulies, Bankim Tejani, Bryan Richardson, and Laura Weiland

Sandia National Laboratories

1 Introduction

Flexible Alternating Current Transmission Systems (FACTS) devices are installed on electric power transmission lines to stabilize and regulate power flow. Power lines protected by FACTS devices can support greater current because anomalies—frequency excursions, voltage drop, phase mismatch, malformed wave shape, power spikes, etc.—that would otherwise cause breakers to trip are removed or greatly reduced by FACTS conditioning.

A FACTS device can also limit the amount of current that flows on a line by effectively increasing the line's impedance. This enables a much greater degree of flow control than provided by a switch or breaker. In particular, when current applied to a FACTS-protected line is greater than the device will allow, the power merely flows elsewhere rather than tripping a breaker, and power continues to flow on the protected line.

Essentially, lines can be run closer to their theoretical capacities when they are protected by FACTS devices. For a large line, that can mean substantial additional power. High-voltage, high-power FACTS devices are building-sized and expensive, but they are lower cost and have less impact per added unit of electric power than new transmission lines. This is the essential benefit of operating standalone FACTS devices on individual lines.

FACTS devices offer an additional benefit: consider an interconnected network where two identical lines are carrying power, one at 50% of its capacity¹, the other at 99%. Assume that any additional load will be supplied equally through the two lines and that there is sufficient generating capacity to support the additional load being considered. Under these conditions, additional load can be supplied only up to the limit of either line, and since one is at 99%, the system can support only about twice the remaining 1% (half of the additional power would go to each line). Additional power would cause the 99% line's protective breakers to trip, at which point all power would attempt to pass through the remaining line, which would then also trip; the generators, being disconnected from their loads, would shut down, and the system would go dark.

However, if the line at 99% were held there by a FACTS device, any added power would go through the 50% line while power continued to flow in the 99% line at its original level. The capacity of this network considered as a whole would be increased by 25%, over and above the stabilizing and regulating benefits provided by the FACTS device. Note that this benefit cannot be recognized by analyzing just the FACTS device and its assigned branch, but only by considering the entire network. For a system that often operates in this sort of unbalanced state, FACTS devices can provide substantial additional capacity simply by forcing more of the network to carry the level of power it was designed to carry.

¹ For this example assume that *capacity* refers to the line's operational limit under local conditions.

This idea leads to a new mode of operation: FACTS Devices can also direct power to less-utilized parts of the transmission network, effectively increasing the capacity of the network, in addition to their customary standalone roles. Because optimum flow for the network as a whole cannot be achieved by considering only single branches, FACTS devices can perform this function only in cooperation with one another. In this report such devices are referred to as Cooperating FACTS devices, or CFDs.

In practice, however, the additional communication required of CFDs opens the potential for subverting the operation of a cooperative system. This report considers both the operational and security aspects of CFDs operating in an electric power system network.

Sections 2 through 7 of this report address the operation of a CFD system²:

- Discussion of the operation of an electric power network (Section 2).
- Description of the elements of a system of CFDs, including existing standalone FACTS device types and their function and operation and the hardware, software and algorithms needed to carry out CFD functions (Section 3).
- Description and discussion of the relationships among the various control and data information flows in a CFD system and discussion of particular challenges posed by obtaining real time data and identifying contingencies in a CFD environment and possible approaches to these (Section 4).
- Discussion of agent management of CFDs including topics related to the environment, roles, interactions, organizations, and operational conditions (Section 5)
- Discussion of the consequences of power system failure (Section 6).
- Discussion of contingencies, how they affect electrical systems and the public, and how they can progress to major failures with the August 14, 2003 blackout as an example (Section 7).

Section 8 addresses security considerations for a CFD environment:

- Assumptions held for this report (Section 8.1)
- Definition of security- and network-related terms and concepts (Section 8.2)
- Description of the Common Criteria methodology for defining generic security policies and identifying specific security functions for devices and processes for a specific system (Section 8.3)
- FACTS Information assets (Section 8.4)
- A description of security features and assets that come from the Supervisory Control And Data Acquisition (SCADA) system used to operate the electric power system (Section 8.5)
- Vulnerabilities of CFD systems (Section 8.6)
- Good security practices (Section 8.7)
- Agent-based security (Section 8.8)
- Discussion of a Network situational awareness and visualization tool developed by Center for Cyber Defenders (CCD) students and its application to analysis of a CFD environment (Section 8.9)
- Security documentation (Section 8.10)

Sections for Conclusions, (Section 9), References (Section 10), and Definitions (Section 11) complete the report.

² As used in this report, a CFD system is an operational group of CFDs; a CFD environment is an electric power system that includes a CFD system as part of its control mechanism.

2 Power Transmission System Reference Model and Context

The Power Transmission System (PTS) is composed of a power network, consisting of generators, transmission lines, and distribution lines; and a control system, consisting of switches, sensors, actuators, etc. Any FACTS devices in use would be considered part of the control system.

The IEEE 118-bus system is the reference PTS for this report. In the analysis of an actual system, determining the extent of the system being analyzed may not be trivial, given that any particular subset of the US transmission/distribution grid is connected to the rest. Usually, the analysis of a system will be limited to the generation and transmission system in the region under consideration, with the distribution systems represented as aggregate loads and some representation of the inertias with other regions. This is consistent with how actual transmission systems are typically operated.

There are several commonly recognized key concerns in operating power systems. The following list of key concerns is abstracted from the final North America Electric Reliability Council (NERC) April 2003 Blackout Report [NERC]:

1. Monitor the power system to ensure thermal limits are not exceeded – Power equipment is designed to operate safely and reliably to a certain maximum level. If that level is exceeded, the equipment starts to overheat, and relay devices isolate the equipment to prevent damage. The system operates so that these limits are not reached and unnecessary outages are prevented.
2. Balance generation and demand – At all times, there must be adequate system generation to meet the loads being delivered to municipalities and commercial customers.
3. Balance reactive power to maintain voltage – In addition to meeting adequate generation, voltages must be kept at appropriate levels because downstream load equipment can operate only within narrow voltage ranges without damage.
4. Keep the system stable – A system could have proper overall system voltage and frequency yet be unstable due to the concentration of power flow in certain regions. Careful control is needed to maintain stable flow.
5. Maintain adequate reserves – At any time, a given generator or transmission line may fail. Additional reserves must be ready to come into service in a timely manner to absorb load when generation failures occur.
6. Plan for normal operations – Adequate planning involves long-term planning to address forecast changes in demand, then increasing generation and transmission assets and maintaining or replacing equipment to meet the projected needs.
7. Plan for emergencies – Emergency plans are sometimes referred to as contingency, remedial, or blackstart plans. They are remedial action plans to address unanticipated conditions and to restart existing or replace lost generation when blackouts and brownouts occur.

These concerns are relevant in the context of this report because a group of CFDs must address these same concerns for the environment in which they are operating. Items 1 – 5 are ongoing daily and moment-by-moment concerns, while 6 and 7 are longer-term concerns.

Item 1 (Monitor the power system to ensure thermal limits are not exceeded), is primarily handled by local relays that detect overload conditions for generators, transmission lines, transformers, etc., and take them out of service before they reach thermal limits. Once a

thermal limit is reached, either the equipment will be significantly damaged or a considerable portion of its life depleted, hence the need for automated relays. Operators should manage the system to avert overload conditions so that equipment does not go out of service because of thermal overloads.

Items 2 (Balance generation and demand), 3 (Balance reactive power to maintain voltage), and 4 (Keep the system stable) are operational concerns, typically handled on a time scale of minutes to hours by a combination of system operators and automated control systems. Some of this automation includes market systems in new deregulated environments in which generation is scheduled and dispatched on an hourly to daily basis depending on the most competitive sources available at the time. Bidding is managed independently by a Regional Transmission Organization (RTO) or Independent System Operator (ISO). Primarily, operators ensure that there is adequate power flow distribution within the system, both in terms of the actual power flows and the reactive flows that affect voltage levels. When low voltages occur, operators activate devices such as capacitors, transformers, and FACTS devices to remedy the system and even out the distribution. Operators also ensure that adequate generation reserves are maintained should a particular unit trip offline, as well as monitoring alarms and line faults in the system. Tools such as state estimators are also employed to predict how the system will react to particular contingencies.

Item 4 (Keep the system stable) is the most complex issue in a power system. This complexity has several relevant aspects.

First, it is difficult to quantify aggregate power system behavior, because of the many interacting components, and to accurately and quickly predict the future state of a power system, because this requires near-complete real-time system state information and high-throughput power flow state estimation software.

Second, the response of ordinary power systems to inputs is nonlinear and affected by many independent elements (resistance and reactance of generators and lines, existing flows, etc.). Under normal conditions, the power system state is either static or changes slowly, which allows relatively easy analysis of system change. However, when a system undergoes significant rapid change due to faults or instability, the equipment parameters also change. Transient and subtransient impedances can be substituted for normal values to simplify analysis of these effects, but these approximations result in unpredictable inaccuracies for multiple interacting devices in a large system.

Finally, there are several classes of stability problems, and each can be further subdivided based on local and regional effects. These classes are interrelated, but the stability problems are different. The following list and information is taken from [Kundur]:

- Rotor Angle (Transient) Stability – The power system consists, in large part, of interconnected, synchronous three-phase devices, primarily generators and motors. Each device has a rotating mass, referred to in general as the *rotor*. When such a device is in steady state, the input mechanical torque and the output electrical torque of the rotor are balanced. When a disturbance occurs, either an acceleration or deceleration of the rotor occurs, causing it to go relatively faster or slower. When this proceeds beyond a certain limit, the affected device loses synchronism and is shut down by protective mechanisms. The concomitant loss of load or generated power causes additional instability. The instability can then propagate to other devices, resulting in a large outage. This has been the most studied problem and it is local to individual machines, so rotor angle stability problems have been addressed extensively, resulting in excitation and damping devices to

control fluctuations and mechanisms for quick identification and isolation of failed machines. For small local disturbances, time frames for the effects are on the order of 10–20 seconds and for large regional disturbances on the order of 3–10 seconds.

- Voltage Stability – Voltage stability is largely concerned with maintaining steady voltage distributions in a power system to keep any one area from having low voltages. Voltage problems occur as a result of either disturbances (e.g., faults, equipment failure, etc.) or changing load patterns over time that cause increases in certain areas that are not properly mitigated. In essence, for a particular line or generator, if there is too much impedance in the system, particularly during heavy loads, the voltage drops in a particular line or generator and leads to low voltages at these locations. Extreme low voltage can lead to brownout conditions can result in power outages if equipment is taken out of service by voltage-detecting relays. Capacitor banks, transformer tap changers, and FACTS devices are used to increase reactance levels to boost voltages in areas where it has decreased and also permits power flows to be maintained. Depending on the nature of the voltage condition and whether it is local or regional, time frames of interest are on the order of seconds to tens of minutes. Longer-term voltage changes can be controlled by operators. Short-term changes, if critical, rely on automation to mitigate. For example, many areas have implemented low-voltage detection and load-shedding schemes that drop loads sequentially to maintain the integrity of the system when voltage drops significantly.
- Frequency Stability – Frequency stability is concerned with maintaining steady frequency within a narrow range of values following disturbances that result in large imbalances between generation and load. Equipment is designed to handle frequency excursions that occur during normal load changes. However, when large changes occur, there may be inadequate generation to cover load, inadequacies in the ability of equipment to respond to the changes, and poor coordination of protection equipment, all of which contribute to frequency instability. These conditions are not modeled in transient or voltage-stability studies because they invoke the actions of processes and controls, such as local device relay settings and inadequate unanticipated device response times, making them the most difficult for which to design mitigations. Often remedies for these problems cannot be discovered until after the fact because they are not modeled. In a large system, if two regions differ significantly in frequency, there can be islanding or isolation of the two systems. Time frames of interest for frequency stability conditions are on the order of seconds to tens of minutes. However, since frequency changes affect equipment controls such as generator outputs, even small changes can have lingering affects for several minutes.

In the development of a CFD environment, each of these concerns must be taken into account. Stability poses the most complex set of challenges. The CFDs will be operating in the context of an existing power system so will be subject to the same constraints and limitations as the existing interconnected power system. Setpoint changes within the CFD environment must be managed in such a way that all of the concerns with balancing power, maintaining voltages, operating equipment below their designed limits, and keeping the system stable are met at all times.

— This page intentionally left blank —

3 Description of a System of Cooperating FACTS Devices

3.1 Standalone FACTS Function and Operations

A FACTS device directly controls the flows on a single transmission line. Indirectly, this affects the flows on other lines within the network by distributing the flows based on the network parameters. FACTS devices use a combination of transformers and Voltage-Sourced Converters (VSCs) to provide control of network flows. Common types of commercial FACTS devices are described below.

3.1.1 Voltage-Sourced Converter

Voltage-sourced Converters (VSCs) are six-pulse converters consisting of six power semiconductor switching devices and anti-parallel diodes. From a direct current (DC) voltage source, the VSC generates a set of controllable three-phase output voltages at the frequency of the system voltage. Pulse width modulation is used to control the firing of the semiconductor switching devices, generating an “average” sine wave. Pulse width modulation also helps mitigate the amount of harmonics.

3.1.2 Static Synchronous Shunt Compensator (STATCOM)

The reactive power exchange between the VSC and the power system can be controlled by varying the amplitude of the VSC output voltage. A VSC that has this capability is called a Static Synchronous Shunt Compensator (STATCOM). If the amplitude of the output voltage from the VSC is increased to above that of the system voltage, the VSC injects reactive power into the system. If the amplitude of the output voltage from the VSC is decreased to below that of the system voltage, the VSC consumes reactive power from the system. Most often, only reactive power is injected or consumed. This is done by keeping the output voltage from the VSC at the same phasor angle as that of the system voltage. A shunt-connected transformer connects the STATCOM to the power system.

3.1.3 Static Synchronous Series Compensator (SSSC)

The magnitude and angle of the injected voltage can be controlled by varying the amplitude and phase of the output voltages produced by a VSC. Such a device is called a Static Synchronous Series Compensator (SSSC). If the phase angle of the injected voltage is kept near the phase angle of the line current, only reactive power can be supplied to the system. However, if the phase angle of the injected voltage has free range, active power can be provided as well. A series-connected transformer connects the SSSC to the power system.

3.1.4 Unified Power Flow Controller (UPFC)

The UPFC consists of the STATCOM and the SSSC interconnected through a common DC bus. This combination called a Unified Power Flow Controller (UPFC) and is significantly more flexible than the separate functions of the STATCOM or SSSC.

With the STATCOM and SSSC connected through a common DC bus, active power needed by the SSSC can be drawn from the shunt-connected converter through the DC bus. This makes it possible for the injected voltage to have any angle with respect to the line current, which in turn gives both real and reactive power control. The functionality of the STATCOM is remains available as well, giving a device with dual functionality and increased flexibility.

In discussion in this report concerning FACTS devices being utilized as cooperating FACTS devices (CFDs), the power electronics and switch level control interactions are assumed to be those of a UPFC. However, other types of FACTS devices can be modeled as CFDs if the limitations of these devices with respect to a UPFC are considered.

3.2 Assumptions about Cooperating FACTS Device Operation

These assumptions make explicit the baseline understanding of FACTS device operation upon which the analysis in this report is based.

1. An operational FACTS device is associated with, and identified by, a branch of a network.
2. Each FACTS device is initialized so that when the device is energized, it is prepared to receive and transmit information necessary to cooperate with other FACTS devices.
3. When a FACTS device does not have sufficient information to adopt a cooperative stance with other FACTS devices, it defaults to predefined setpoint levels or specific operational protocols according to the set policies.
4. The normative behavior of a standalone FACTS device is the regulation of power flow on its assigned branch according to a set of parameter values obtained independently of the FACTS system (e.g., predetermined settings based on policy and/or economic considerations).
5. The essential function of a CFD (as distinct from the function of a standalone FACTS device) is to maintain a higher or lower current flow on its assigned branch than if no cooperation were occurring.
6. The power flow to be maintained by a given CFD is established in cooperation with the other CFDs within the overall CFD environment.
7. CFDs cooperate with one another to achieve and maintain a network state specified by a designated policy. The policies involve managing CFD setpoints according to the flow-allocation algorithm in use and following rules for interactions.
8. CFDs communicate with one another by some means, such as the IP protocol discussed in Section 8 below.

3.3 Hardware/Software Integrated System

3.3.1 Algorithms of Interest

This section describes the algorithms necessary for a CFD system that would not be needed for a power system with no CFDs.

- Initialization
 - At system startup the hardware and software elements of the system are initialized according to a specified process. When an individual branch controlled by a CFD is de-energized and then re-energized, the settings for that branch need to be determined according to a specified process. Settings based on network conditions would involve communication with other CFDs and/or the network's SCADA system.
 - The system must be in a steady state condition before long-term control can be activated. The long-term control process must be notified when the initial steady state has been reached. Dynamic control might begin at initialization using limits based on policy or on the limits of the individual branch, or it might wait for the establishment of steady state and information from other CFDs. A two-stage process is possible in which policy-based settings are replaced by cooperatively derived limits. How this occurs is important to the operation of the system but is not addressed in this report.
- Dynamic Control – Dynamic control responds to power fluctuations on a transmission line and provides control when transitioning between FACTS setpoints. When multiple CFDs are operating, it is important that they adopt and put into operation new setpoints in such a way that significant transients and instability do not occur.

- Long-Term Control³ – Long-term control uses a flow-allocation algorithm to identify a desired network configuration and the branch setpoints that the configuration implies, in particular a set of setpoints for the branches controlled by FACTS devices. Setpoints of other branches in the network may also be specified for monitoring purposes, but they cannot be directly controlled by any system element⁴. The values of the FACTS-controlled setpoints are derived from the optimal power flow on each branch of the network. Branches controlled by FACTS devices can then be constrained to allow no more than the amount of power determined by the algorithm. Determination of which branches in a network should be controlled by FACTS devices is the subject of research; see, for example, [Chaloupek], which discusses the use of genetic algorithms to determine effective FACTS device placement.

The Ford-Fulkerson optimization algorithm has been described in use for the FACTS case in [Armbruster1], [Armbruster2], [Armbruster3], and [McMillin]. This algorithm enables the maximum power flow through the network. Another possibility, of interest because networks are often not at capacity, is to minimize the variance in the fraction of branch capacity being carried by each branch over all system branches. When measured by this metric a network in which some branches were carrying very little power and others were nearly at capacity would not score as well as one where every branch was carrying nearly the same percentage of its capacity. This metric would allow a result comparable to a maximum flow condition if a value near 100% could be achieved. Gradient-descent-type algorithms have also been considered.

[Vlachogiannis] presents an ant colony system (ACS) method for network-constrained optimization problems in which the constrained load flow (CLF) problem is a distributed combinatorial optimization problem. Cooperating artificial “ants” cooperate to find an optimum solution to the CLF problem. A pheromone matrix in the role of global memory provides the cooperative framework. The ACS algorithm is applied to the IEEE 14-bus system and the IEEE 136-bus system. The results of the ant-based algorithm, a probabilistic CLF algorithm, and reinforcement learning (RL) methods are compared to show the benefits and flexibility of the ACS algorithm.

Each algorithm has strengths and weaknesses in terms of the computational requirements, accuracy, communication, and processing required of the CFDs in an operational environment. From a security perspective the relevant concerns are the amount of communication needed to execute the cooperative result and the effect of incorrect or missing information on the distributed computation. The more communication that must occur, the greater the exposure of the computation to propagation of errors and the greater the effect of perturbed information flow, the greater the consequence of a disturbance. There is inherently more risk in higher-communication higher-consequence computations.

3.3.2 Hardware Elements

This section discusses the hardware needed in addition to the transformers and VSCs discussed in Section 3.1 to enable FACTS device functionality.

³ “Long-term control” is control of phenomena that occur on a scale of seconds/minutes. The distinction is needed because FACTS devices also provide dynamic control, which deals with phenomena that occur at time scales of micro- to milliseconds.

⁴ They may be *indirectly* controlled, however, by the CFD system as a whole.

- DSP Board – The Digital Signal Processor (DSP) board obtains analog information (voltage, current, etc.) from the power electronics sensors and processes it for the controllers, and contains the IGBT controls that enable device power electronics to execute the setpoint changes. The DSP control elements are typically hardware encoded, since their operations are deterministic.
- Power Electronics – The power electronics consist of the voltage and current sensors, transformers, inverters, capacitors, etc., that interface with the power line to carry out the setpoint changes determined by the long term and dynamic controls. They are mediated by the DSP. As with the DSP, these functions are hardware encoded.
- Local Power Line – Though this is not part of FACTS device itself, it is important to include given that the flows on the local power line are sensed and controlled directly by the FACTS device. At the University of Missouri/Rolla, a hardware-in-the-loop simulator that includes resistive loads and a motor-driven generator is used to produce actual local power line flows that are sensed and modulated by the CFD devices.

3.3.3 Software Elements

- Embedded Computers – Dynamic and long-term control is implemented in software residing in an embedded computer within each FACTS device. In FACTS devices currently operating, this computer has a commercial off-the-shelf operating system (OS). This is not likely to change as FACTS devices are more widely deployed. The operating system is the source of significant vulnerability to everyday hacker intrusions, viruses, worms, trojans, and other “malware.” This is primarily because the common operating systems are widely exposed to adversarial experimentation and are attractive targets because of their widespread use. Conversely, hardening and security maintenance of the OS at the system-administrator level (maintaining accurate access control lists, eliminating unused processes, consistently upgrading and patching, active firewall, etc.) is the single most relevant act in protecting a standalone FACTS installation. Communications packages, agents, and algorithms that support FACTS functionality for both standalone and cooperative operation should be treated as specific applications and as such are subject to maintenance updates and are themselves potential targets of adversarial activity.
- Digital Signal Processor – The onboard processor of the FACTS device’s digital signal processing (DSP) board is considered part of the FACTS device. The operating system and software running on the DSP are subject to maintenance updates and present another security risk. This computer will probably not utilize a conventional operating system since it is not a general-purpose computer, although reduced-footprint Linux systems are increasingly common in such applications. Although somewhat less vulnerable than the device’s general processor, the DSP is more difficult to harden and maintain since it customarily has a more esoteric OS and interface. Nevertheless, successful adversarial penetration of this processor can result in complete malfunction of the FACTS device since this is the route by which information about the electric power system enters.
- Switch-level control – Software running on processors at levels near the hardware switching elements is likely to reside in electronically programmable read-only memory (EPROM) or other near-hardwired devices. It is less exposed to the outside network and harder to modify. It should be included in a risk analysis, however, because of the potential for complete cooption of device function if penetration occurs.

Figure 1 is a representation of how information flows occur in an individual FACTS device for the hardware and software elements described above.

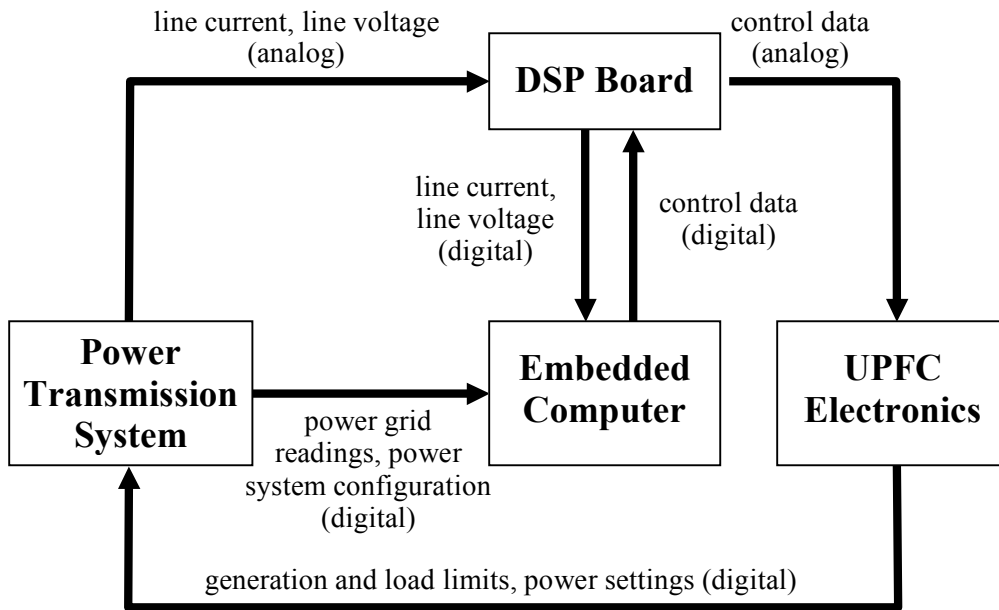


Figure 1. Information flows in and around an individual FACTS device

3.4 Cooperating FACTS Function and Operations

A single Flexible A/C Transmission Systems (FACTS) device directly affects flow on its assigned branch. In changing the flow on its own branch, however, the CFD also indirectly affects flow on the other branches in the network, the nature and magnitude of which can be calculated using Kirchoff's laws. The combined direct and indirect effects of multiple well-placed, correctly operated FACTS devices throughout a network can result in near-optimal power flow on a network-wide basis.

This section describes the baseline desired behavior of a set of CFDs providing long-term control of a network. The CFDs establish a set of setpoints, one for each of the branches controlled by the individual CFDs. In practice, recalculation of setpoints can be driven by any of several causes: a predetermined schedule, a loss of a branch or generator, a large increase in load, and other events⁵ that significantly alter the power system's load carrying requirements. This cooperative behavior is referred to as "long-term control" in the High Order Object-oriented Modeling Technique (HOOMT) diagrams [Ryan]. Setpoints based on such cooperative goals are referred to as long-term setpoints.

A FACTS device is made capable of cooperating by adding algorithms and communication capability to the standalone device. These allow it to communicate with other FACTS devices and to receive information about parts of the power network of which it would otherwise remain ignorant. The added communication capability enables execution of distributed flow-allocation algorithms such as distributed MaxFlow.

⁵ Events trigger computation of long-term setpoints either by placing the network into an undesirable state or changing the world state to the extent that the long-term setpoints in force may no longer be functional.

3.5 Computation of long-term setpoints

The outcome of any computation intended to provide appropriate power system control information to a group of CFDs is a set of flow limits, one for each branch controlled by a FACTS device⁶. These flow limits are referred to as setpoints. For the purposes of this report, algorithms capable of providing outcomes of this sort are called flow-allocation algorithms. The general MaxFlow algorithm [Armbruster2] is one such algorithm.

Computation of long-term setpoints proceeds approximately as follows:

1. The CFDs begin with a mutually agreed-upon current network state as an initial reference. The current state can be used to back track in the case of a failure or damage that alters the network. Depending on the details of the algorithm to be executed, the network state may be for the global network or just for a portion of the network that falls under each CFD jurisdiction.
2. The CFDs arrive at a mutually agreed-upon desired network state.
 - a. This may be established *a priori* by policy, according to the selected flow allocation algorithm; for MaxFlow, the mutually agreed-upon, desired network state is any state in which no branch of the network exceeds the maximum flow state.
 - b. According to the HOOMT diagrams in [Ryan], long-term control uses the current network state, the `compute_next_level_setpoint` method, and FACTS power system configuration information to calculate the setpoints that give maximum flow.
 - c. Observation suggests that optimization is not always necessary (e.g. when the amount of power flowing in the network is significantly less than its capacity)⁷. This implies that the network state in such cases is within the parameters of the desired network state. This is the case in the practical MaxFlow algorithm, where the network's operational steady state is the desired state and, as a result, no change is perceived to be necessary. When no contingencies or major state changes occur, the setpoints don't change, and each CFD simply maintains its setpoint.
 - d. It might be the case that the capacity of a well-designed network under ordinary conditions of load and generation will never be exceeded; that is, the network would be in an undesirable state only in the case of contingencies that significantly alter the network's ability to produce, deliver, or use power.
3. The CFDs produce a mutually agreed-upon set of discrepancies between the desired and current network states. The discrepancies may be predefined as variances from some baseline condition and left for the individual devices to deal with. In the distributed case, information about the discrepancies is communicated to other CFDs, as in distributed MaxFlow. Section 7, Contingency Analysis, discusses how contingencies cause these discrepancies to occur.

⁶ It is not necessary for the set ever to be assembled in one place, but each CFD must have the flow limits for its branch as a result of the computation.

⁷ The *capacity* of a network is the placement and magnitude of all loads when the network is in a state such that additional load anywhere in the network causes breakers to trip. This is the "maximum flow" referred to in [Armbruster2]. Note a network being at capacity does not necessarily mean that every individual branch is at capacity.

The primary undesirable condition is an *unbalanced* network, i.e., some branches are carrying more than average loading⁸ (with greater line loss) and some are carrying less. If the network as a whole is under capacity, this is not of great concern; efficiency might be lower than desired but network operation is not in jeopardy.

If the network is at capacity, the unbalanced condition implies that some individual branches are under capacity. Although the under-capacity branches could carry more power, some fraction of any added power will flow through branches already at capacity (by Kirchoff's Current Law) and cause them to overload. In essence, there is unused⁹ transmission capacity in a network that's at capacity but unbalanced.

Note the limiting capacity needs to be carefully considered for each electrical device. For example, power lines have different thermal, load, and stability limits depending on the length of the line and the configuration of the rest of the system, so the limiting factor and capacity condition will depend on which of these limits is most constraining for each line in the system. Similar caution applies to generators and other power equipment.

4. CFDs produce a set of setpoints intended to reduce or eliminate the discrepancies between the current network state and the desired network state.

The major benefit of utilizing CFDs is that branches at or near capacity can be constrained to carry no more power than their maximum capacity permits, i.e. held at 100% of allowable loading. Under such conditions, additional power (again based on KCL) would seek some other path (i.e., less-loaded branches). Note that if branches exceed short term overloading capacity for too long, they will open up via local relays. CFDs must be able to respond to avoid as many of these conditions as possible.

5. Each CFD acts to maintain the setpoint for its assigned branch.
 - a. The setpoint for each particular branch is sent to dynamic control where it is used along with local branch data (from current and voltage sensors that sit within the UPFC device) to calculate the setpoints needed to transition the branch from the current state to the desired state using the `compute_next_level_setpoint()` method (see [Ryan]).
 - b. Information would also be sent here in the wake of a contingency, requesting dynamic control's help in mitigating it.
 - c. This control data are then sent to the DSP board. Here, IGBT control takes the control data, along with local branch data (again, from the voltage and current sensors that sit within the UPFC device), and calculates the switch settings for the voltage source inverter using the `compute_next_level_setpoint()` method.
 - d. From here, the switch settings are sent to the IGBT Driver and Protection Board within the voltage source inverter. The driver then sends `switch_on/switch_off` commands to each of the six IGBT switches according to the switch settings that are sent to it which in turn with other power electronics elements, maintains or changes setpoints.

⁸ "Loading" is an imprecise term used here to refer to the amount of power a branch is carrying as compared to the amount it is able to carry (i.e., its capacity).

⁹ Such unused capacity is not merely unused; without cooperating FACTS devices, it's unusable.

— This page intentionally left blank —

4 FACTS Information Flows & Repositories

Figure 2 is referred to in the following discussion of the information flows that occur in carrying out the operation of cooperating FACTS devices (CFDs) in an electric power system. The illustration depicts the basic interactions of CFDs in a power system. The interactions occur in the context of a power system with generation, transmission, distribution, and end users, and each FACTS device controls the flows in a particular transmission line.

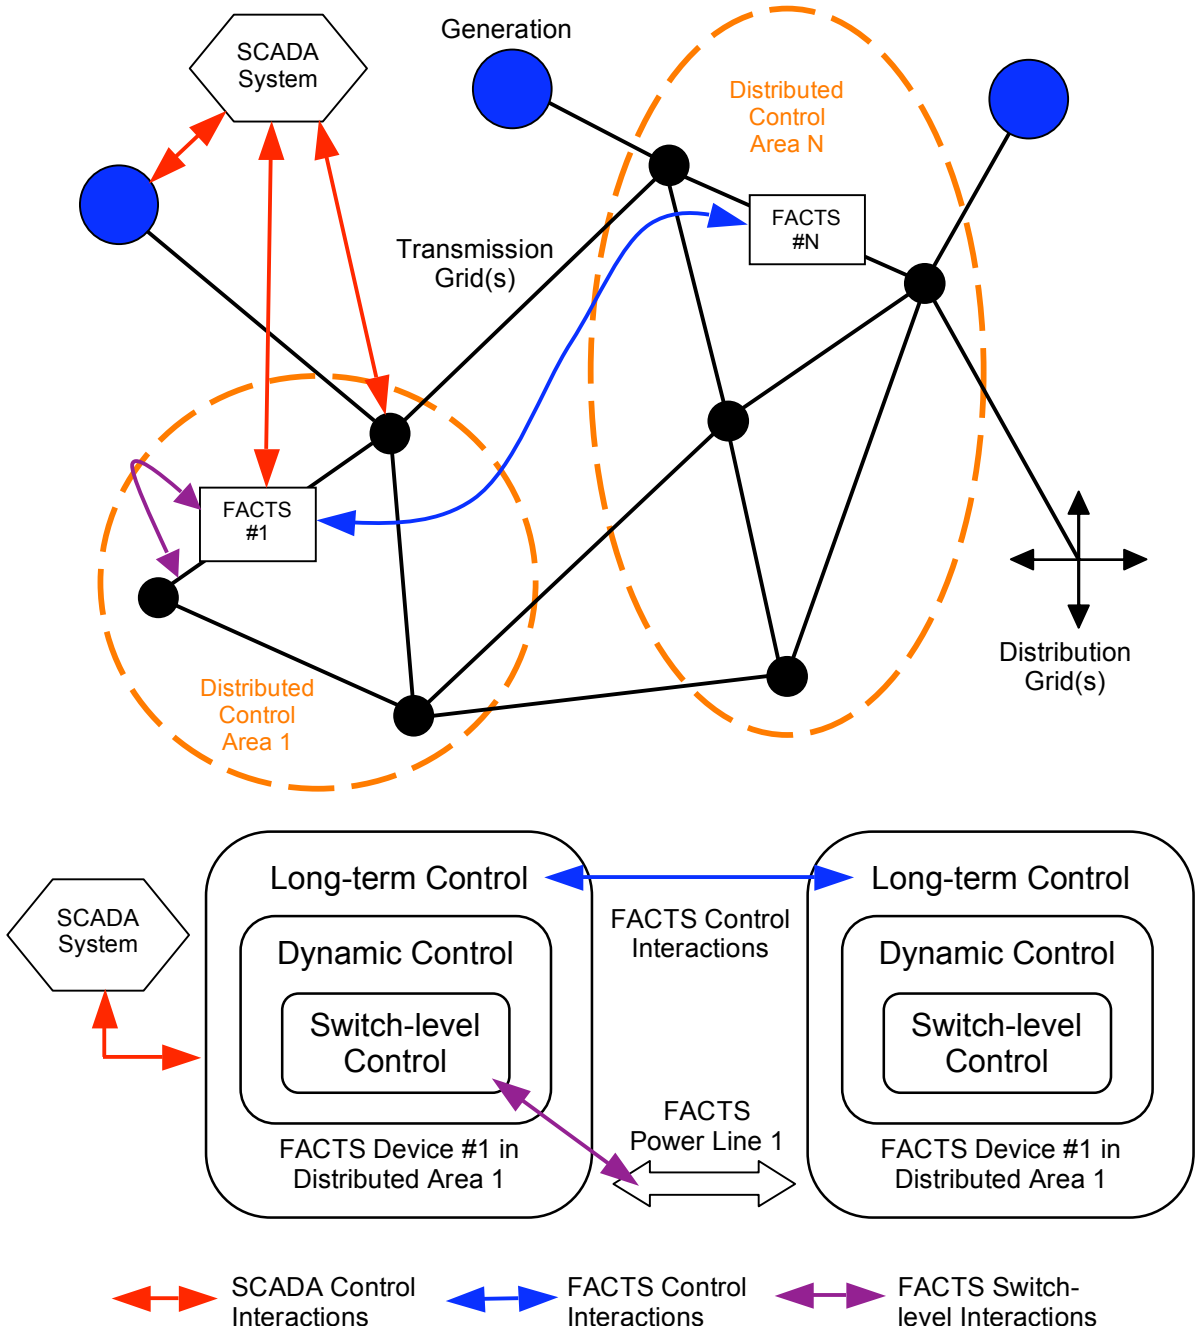


Figure 2. Notional FACTS layout and its relationship to the power system

As with a standalone FACTS, a CFD system will interact with the power system controls via the Supervisory Control and Data Acquisition (SCADA) system to maintain setpoints on individual branches.

The long-term control interactions involve both the interactions necessary for the execution of the optimization algorithm (e.g., Max Flow) computed collectively by the FACTS devices and all CFD-related ancillary and management communication, such as managing execution of setting changes and verifying that other FACTS devices are valid.

Some interactions occur only in a CFD environment. Interactions between the CFD System and the SCADA system will include the same control, setting, and status information relayed back and forth between the SCADA control system and the individual FACTS devices. In addition, each FACTS device will obtain system status data from the control area that it manages in order to know the current network state and compute the desired state as described above. Specific problems that may occur with these data needs (error correction, synchronous data, and limits in data availability) are addressed in Section 4.6.

Dynamic control must avoid interaction among FACTS devices that can generate transients in the system when executing the long-term control settings. Each of these interactions needs to be considered independently and in relation to the others.

4.1 Switch-Level Control Interactions

Switch-level control interaction involves the sensor information needed by the FACTS device for monitoring and the executed control actions that are based on information passed from dynamic control to the DSP, which determines the reference signals that set the value of the branch power flows. Once the reference points are set, onboard power electronics execute the setting by injecting current and voltage into the power line through the inverters and transformers to modify the power flow to the set values. The switch-level control interactions of a CFD are the same as those of a standalone FACTS device. The differences lie in how the long-term and dynamic controls work to modify the setpoints that are passed on to switch-level controls. Figure 3 shows the arrangement of switches in the UPFC.

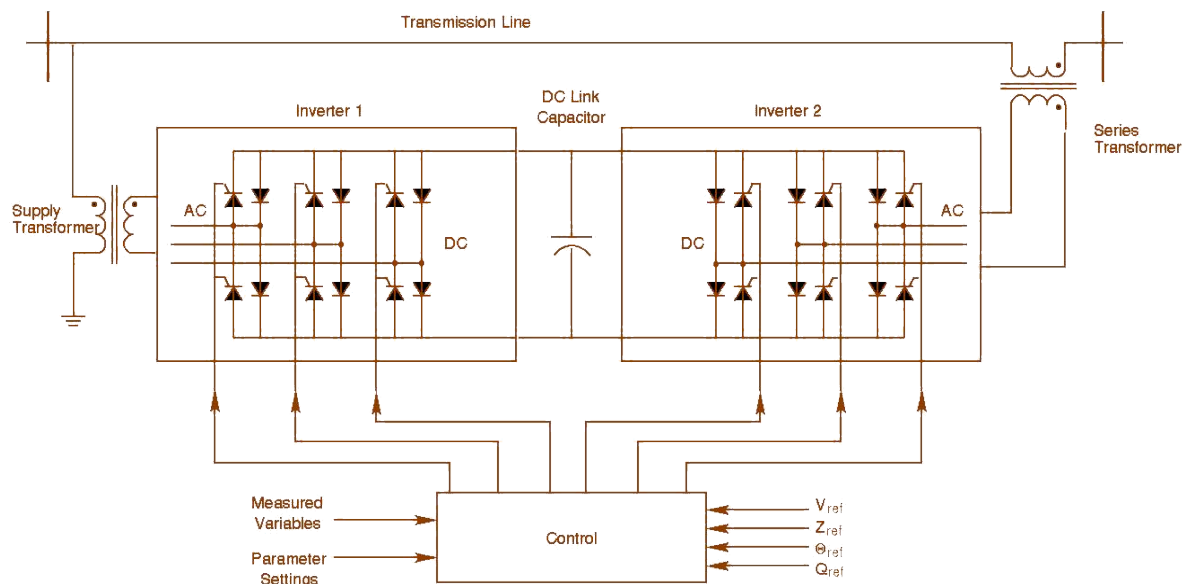


Figure 3. Switch-level architecture of the UPFC FACTS device

Table 1 shows the types of Input/Output (I/O) utilized by the switch level control interactions in a CFD modeled as a UPFC FACTS device in a CFD environment.

Table 1. Switch-Level Control Interactions I/O

Inputs	Type	Source	Purpose
Local Readings (Voltage, Current, Power)	Analog	Power devices (CTs, PTs, etc.)	Gain local power line information
Reference Settings (Vref, Zref, Qref, θ_{ref})	Digital	Obtained from Dynamic and Long Term Control	References to determine how to apply settings the CFD line
Parameter Settings and Measured Variables	Digital	Derived from combination of fixed settings, current operational setpoints, and system conditions	Ensure that the CFD operates within both internal limitations and system constraints

Outputs	Type	Source	Purpose
Inject Voltage and Current	Analog	Series and Shunt Transformers	Execute power flow changes based on setpoints
Switches	N/A	High voltage interruption devices	Place all or part of CFD in or out of service

4.2 Dynamic Control Interactions

The dynamic control in a CFD works to ensure that when settings are changed, two different CFDs won't try to execute changes at the same time or a change made in one CFD doesn't trigger a responsive setpoint change in other CFDs. Either of these cases can cause stability problems. Dynamic control moderates the implementation of changes dictated by long-term control. The only interaction between dynamic and long-term control is the transmission of setpoints from long-term control to dynamic control (Figure 24 in [Ryan]).

4.3 Long-Term Control Interactions

Figure 4 depicts the information flow surrounding the maximum flow algorithm. Current system state and contingency information are passed to long-term control from the PTS through the DSP board or from the local SCADA system. When the optimization algorithm finishes, the desired network state is then sent to dynamic control for execution.

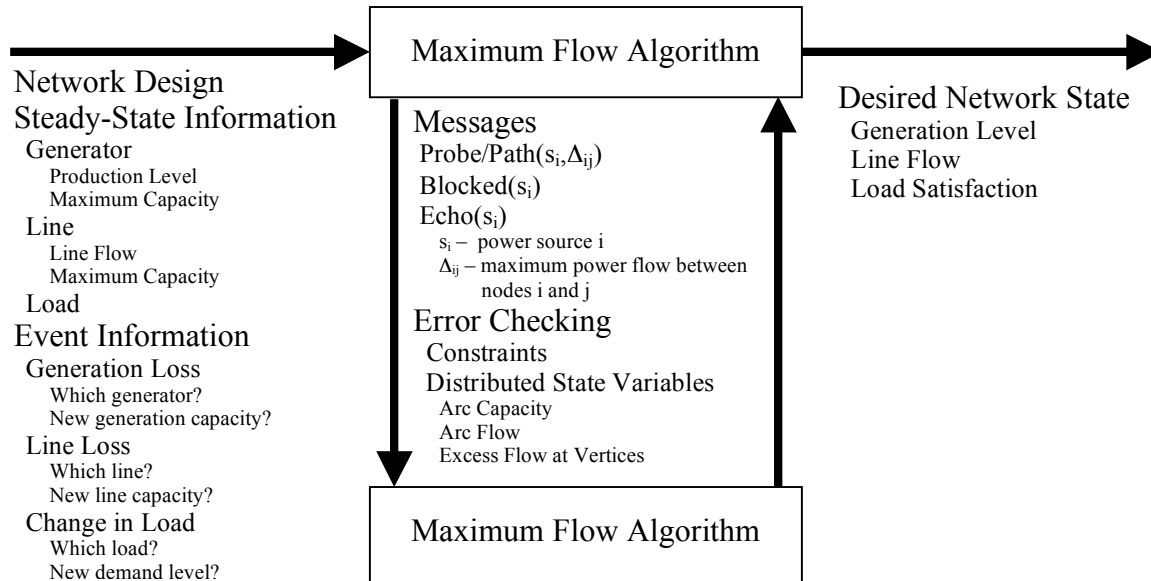


Figure 4. Maximum Flow Algorithm Information Flows

Information is passed between FACTS devices when executing a distributed maximum flow algorithm as shown. Messages are passed between the long-term control algorithm instantiations when individual flow paths cross regional boundaries [Armbruster2], [McMillin]. Distributed state variables such as arc capacity, arc flow, and excess flow at the vertices can be included in the communication as a form of error detection in the distributed version of the algorithm [Armbruster3]. Table 2 shows the types of I/O utilized by the long-term control interactions in CFDs in a CFD environment.

Table 2. Long-Term Control Interactions I/O

Inputs	Type	Source	Purpose
Dynamic Control Feedback	Digital	Local CFD Dynamic Control	Obtain modifications to setpoint change implementation to prevent oscillations
Data Exchange with CFD neighbors	Analog and Digital (Ethernet)	Neighbor CFD	Data necessary to implement distributed max flow algorithm
Control Exchange with CFD neighbors	Digital (Ethernet)	Neighbor CFD	Information necessary for cooperative agreement on CFD changes
Outputs	Type	Source	Purpose
Dynamic Control Feedback	Digital	Local CFD Dynamic Control	Pass computed changes and next setpoints to Dynamic control to execute setpoint changes
Data Exchange with CFD neighbors	Analog and Digital (Ethernet)	Neighbor CFD	Data necessary to implement distributed max flow algorithm
Control Exchange with CFD neighbors	Digital (Ethernet)	Neighbor CFD	Information for cooperative agreement on CFD changes

4.4 Software Code Elements

Software code is located throughout the FACTS device to carry out the programming necessary for CFD operation. The majority of the code, for long-term and dynamic control, resides in the embedded computer. The DSP board and UPFC power electronics code also execute CFD-specific code. Different types of access to this code are required. For example, different levels of user privilege may be required to read system settings and sensor data, to change settings, and to change program files. This would require different user privileges to be defined and safeguards for protecting the information.

4.5 SCADA Interactions

SCADA interactions involve all of the information exchanged between the SCADA system and the FACTS device. In a standalone system, as noted above, the information will primarily be the readings from the FACTS branch, and possibly control signals which allow an operator to change the settings of the FACTS device remotely.

In the setting of a power system with CFDs, the FACTS devices obtain power system network state information. As defined previously, the network state is the composite set of flows for each branch in the power system and incorporates the underlying network topology (impedances), equipment states (in-service, out), and voltage, current, and phase angles for system branches and nodes, which collectively indicate the overall state of the system.

The amount of network state information required at each CFD will depend on several factors such as the optimization algorithm utilized, the subset of nodes and branches of the entire network needed to do an adequate optimization. Table 3 shows the types of I/O utilized in the SCADA interactions by CFDs.

Table 3. SCADA Interactions I/O

Inputs	Type	Source	Purpose
SCADA Analog System Status Data (voltage, current, frequency, etc.)	Analog (via Serial or Ethernet links)	Control Center (aggregated from field devices in remote stations)	System status data used by CFDs to compute distributed max flow and determine setpoints
SCADA Digital System Status Data (switch positions, equipment on/off)	Digital (via Serial or Ethernet links)	Control Center (aggregated from field devices in remote stations)	System status data used by CFDs to sense contingency changes to help determine setpoints
CFD Control Settings	Digital (via Serial or Ethernet links)	Control Center	Control settings to remotely program and override CFD device settings (if applicable)
Outputs	Type	Source	Purpose
CFD Control Settings	Digital (via Serial or Ethernet links)	CFD	Inform Control Center of CFD Settings and current setpoints
CFD Analog and Digital Power Line Data	Analog and Digital (via Serial or Ethernet links)	CFD	Inform Control Center of CFD line readings and status (in-service, out of service)

4.6 Special Challenges in a CFD Environment

Several special challenges are posed by the new types of information and operational considerations required by a CFD environment—challenges related to identifying contingencies and those related to how to process incomplete information obtained from the SCADA system. These new challenges are posed by the introduction of the use of CFDs to maximize power flow in power systems. Of course, automated systems such as Automatic Generation Control (AGC) and Remedial Action Schemes (RAS), as well as others, are used effectively in existing power systems. However, these systems respond and control individual devices to specific local settings and do not attempt to coordinate power flow for a region. The development and implementation of these systems has proceeded through a long history based on simulation and modeling, planning, and experience. In existing power systems, response to regional power flow change is effected through human operators, with the help of specified procedures and automated systems such as AGC.

The uniqueness of a CFD environment is the attempt to control power flow regionally as well as for each defined CFD line. It is anticipated that the same combination of simulation, modeling, and experience will be necessary to properly address these challenges should a CF environment be utilized in future power systems. These specific challenges are discussed below, with a few comments on how these may begin to be addressed.

4.6.1 Recognizing and Responding to Contingencies

In developing a CFD system to identify, recognize, and respond to contingencies, the goals are to “first, do no harm” by maintaining a level of reliability at least as good as that of the non-CFD system, and to optimize the power flow.

It is supposed that the distributed flow allocation algorithm used by a CFD system (nominally the MaxFlow algorithm, but see Section 3.3.1, Algorithms of Interest) to compute setpoints for the CFDs will be executed on a regular, ongoing basis. Based on the results of the distributed computation, the CFDs will agree on the necessary setpoint changes, negotiate the order for executing them, and execute them. The CFDs will need to be able to recognize and respond to contingencies that occur in a system in at least two ways:

First, when a contingency occurs, the system state will change, which means the setpoints for the CFDs might need to be re-computed. Thus the distributed computation should be executed to respond to contingencies that change the system state.

Second, if a contingency occurs in the process of negotiating changes, there must be safeguards, such as suspending a new change or specifying a process for dealing with these conditions, to prevent these changes from amplifying the stability problems resulting from the contingency. An approach to address these issues is to simulate and model what happens in a CFD environment in response to specified contingencies and during the execution of changes to the system and make adjustments to policies and processes based on the results.

The trigger for re-computation of the distributed allocation algorithm depends on how the contingencies of interest are recognized. Section 7 discusses the kinds of contingency that can occur, including transmission and generation equipment failures and large load changes. Relevant questions in this context are: Which of these classes should trigger recomputation and what signatures should be used as triggers?

If the system recognizes and responds to too many small contingencies, the constant flood of change requests will prevent normal operation of the CFDs. On the other hand, if the definition of contingency is too restrictive, the system may not respond to important events. One approach is to define a set of contingencies (such as the pre-identified N-1 contingencies identified by system studies) to which the system will respond and re-compute the allocation algorithm whenever such contingencies occur. From there, other types of situations, such as heavily loaded lines or low-voltage conditions, can be considered as triggers. Once the CFD system is responding to a set of contingencies, operators can begin collecting performance data and honing the trigger conditions based on statistical analysis of the system's behavior.

An equally important question is how to differentiate between true contingencies and false triggers generated by an adversary. The general approach is to prevent malicious attacks of this type, but if this fails the system may be placed in a denial-of-service flood of false contingency triggers or be misled into recomputing false results by bogus data. As above, a data collection/statistical analysis result might yield distinguishing characteristics, but, as is always the case with malicious behavior, it's the one that the system hasn't learned to recognize that causes problems.

4.6.2 Data Completeness

A CFD power system environment uses an allocation algorithm (e.g., MaxFlow) to calculate the settings for the CFDs. Accurate, near-complete information about appropriate system component states is needed to obtain optimal results. This section discusses this challenge and how it might be addressed.

In existing power systems, there may not be sufficient information available in real time to completely describe the system, for several reasons. Not all devices have sensors; noncritical equipment may not be monitored at all. Older or non-standard sensors may provide only partial information and malfunctioning sensors can give incorrect information. In existing SCADA systems, information is retrieved from remote stations according to a predefined polling sequence and can be anywhere from two to ten seconds old. This means that there will be a lack of synchronism for the retrieved information.

In existing power systems, there are no controls that depend directly on real-time power flow information, so a lack of information usually does not present a problem in operating power systems. On the other hand, the calculation of distributed max flow, for instance, depends on

complete information being available. A deployed CFD system should incorporate methods to counteract the inevitable lack of complete information.

In a CFD environment, each CFD will have an area of control from which to get its portion of system data (see Figure 2). Each area will most likely overlap adjacent areas. To compute max flow, or any similar algorithm, there must be sufficient information about the system buses and branches to allow computation of balanced power flows. The areas must be chosen in such a way that the inputs and outputs to the system balance and there must be adequate information about the system to enable successful completion of the distributed calculation.

In a SCADA-based environment, each CFD will obtain system information from the power system control center. Dedicated sensor and data systems could be implemented for the CFD system, but this would probably be prohibitively expensive. In existing SCADA systems, data from remote sensors arrive at the control center after a 2- to 10-second delay because the remote sites are cyclically polled from the control center. We assume all such data is time-tagged, but data available for a given calculation might not all refer to the same time period¹⁰, even if it is complete per the previous paragraph. If the system were changing rapidly, the magnitude of differences in values and asynchronicity could preclude completing the power flow calculations, and thus, for CFDs, determining maximum flow. This doesn't affect existing systems, because they don't require real-time data for executing controls¹¹.

Any approach to this challenge should include examination of the power system state estimators used in existing power systems. Power flow computation in a CFD environment is analogous to state estimation in a conventional SCADA system; in both cases, input is incomplete and out of sync and the results are used in making control decisions. In conventional systems, however, results are not used for direct control, only as input for human decisionmakers determining how to respond to contingencies and allocate resources.

State estimators analytically estimate the flow on each branch of a network topology utilizing available system data (line impedances, voltages, power flows, equipment availability, etc.). They are essentially power flow calculators modified to execute with incomplete information, hence their relevance to CFD computations. In general, these methods determine the consistent system power flow state that gives the lowest level of mismatch between results and observations and report this as the most likely current system condition. This is adequate in conventional systems since the estimated system state does not directly control the system.

An approach to developing a CFD optimization algorithm based on incomplete information would be to compare the results of running a given algorithm on two different data sets, one complete and the other from a state estimator using asynchronous and incomplete information. If the differences are great, it might be possible to develop a CFD algorithm that can accommodate the information limitations inherent in existing systems. If the errors caused by having incomplete information were such that setpoints computed by a CFD system with incomplete information were very close to those computed by a CFD system with perfect information, then that algorithm could be effectively utilized in existing systems with these information limitations.

¹⁰ At least, there may be no designed system mechanism to produce temporally consistent data.

¹¹ There are exceptions, but they are usually limited in scope; e.g., power line protective relays and other automated relay protection mechanisms use specific local data taken in real time.

— This page intentionally left blank —

5 Using Agents to manage FACTS devices

An agent theory specifies what an agent is and what it does. This section discusses the application of agent technology to power management and outlines an agent theory for operating a system of Cooperating FACTS Devices (CFDs). The agents of interest are referred to as *CFD agents* in this section. This section should be considered a set of guidelines for designing a CFD agent system. Additional comments on agent technology specifically related to security appear in Section 8; See Section 8.8.4, Agent-Based Mitigation Strategies, and Section 8.10.1.3, Agent-Based Security Policy.

We recommend, and follow in this section, the Gaia process [Zambonelli], which considers the relevant issues in pragmatic order to produce the representational elements that make an agent theory a good agent theory: What the system knows about the world; how its actions are guided; the system components and their relationships with one another; how the system's information state changes over time; and how the environment affects the system's information state. [Zambonelli] is especially useful because it provides as examples application of the process to two rather different use cases. The authors also recommend [Rehtanz], an extensive body of work specifically tying agent technology to electric power control and operation.

5.1 Consideration of Agent technology for Electric Power Management

Whether to apply a particular technological approach should be based on whether the application requires the benefits offered by the technology¹². Power system resources need to be operated and managed by a distributed system that includes at least several entities who¹³ allocate resources, negotiate trades, share workload, provide redundancy, and maintain cybersecurity through mutual observation and response. Such elements would need to be social; the real-world system could not be operated by elements that could not communicate with one another. These elements should be autonomous, because the rapidity of power system phenomena dictates the ability to act with authority without seeking approval in real time. Finally, they should be situated, because the power system is made up of electromechanical components that require intervention and for which the correct control inputs can be ascertained only by observation of device state.

Agents have all three capabilities. An *agent* is a software process (or set of processes) that is autonomous (can act independently of outside influence), situated (receives information from and acts upon its environment), and social¹⁴ (communicates with and forms organizations with other entities). Available agent frameworks possess these properties to different degrees, implement them differently, and may have other properties of interest, but these three are approximately definitive. Software missing one or the other of these properties might be called an agent, but software with all three could hardly be called anything else. In the end, the term “agent” is shorthand to indicate that software has (or needs for design purposes) these properties. An important component of the CFD agent theory is that the means by which security conditions and responses to security policy violations are handled is policy enforcement by closed coalitions of “good” agents.

¹² At least in part; cost and cost-benefit ratio need to be addressed as well, but the benefits are usually at least notionally developed before the more-difficult subject of cost is raised.

¹³ We might have said “that” instead of “who”, but at present most of these “entities” are human.

¹⁴ Being social is a specialization of being situated, but qualitatively different in that it requires different principles, protocols, languages, and ontologies.

CFD agents would almost certainly be deployed initially as a closed system; that is, a homogeneous system providing standardized services designed around a single species of agent. Individual agents of this closed system would be innately cooperative and trustworthy and have goals based on policy.

Movement into an open system is fundamentally necessary in a market-based system, so that self-interested profit-motivated agents of different species can interact. Open systems are more difficult to design because the set of possibilities for interaction, reaction, and motivation is broader than in the closed case. In an open system, agents using different ontologies, algorithms, and motivations interact with one another to trade and transfer power. In addition, it is likely that systems of unknown provenance will interact with the power system agents in unspecified ways; an open environment essentially forces security concerns.

5.2 Environment

Gaia analysis of an agent-based system requires specifying the environment, the roles of the agents involved, how they interact with one another, how they are organized, the rules required of the organizational members, and the liveness and safety conditions to be maintained by the agents.

The *environment* of an agent system consists of all the things external to the system agents with which they must interact to fulfill their roles. We recommend defining the environment as a class hierarchy, which can then be used as part of the system ontology by the agents.

The primary classes needed by agents dealing with a CFD system:

- I. Electric power system component
 - a. Source
 - b. Load
 - c. FACTS device
 - d. Branch
 - e. Bus
 - f. Switch (breaker)
- II. Entity
 - a. Human
 - b. Agent
 - c. Organization
- III. Power Flow

5.3 Roles

A *role* is a group of related goals and functions. CFD agents have five roles:

- I. Human interaction
 - a. Act on executable commands
 - b. Respond to requests for information
- II. Power system interaction
 - a. Normal operation
 - b. Contingency response
 - i. Contingency recognition
 - ii. Fault recovery

- III. System state prediction, planning, and goal derivation
 - a. Long-term power flow calculation (MaxFlow, Gradient descent, etc.)
 - b. State prediction based on statistics
- IV. Information management
 - a. Network state marshalling
 - b. Common operating picture maintenance
- V. Policy Enforcement (see section 8.10.1.1, “Security Policy Elements”)
 - a. Maintain safety conditions engendered by policy
 - b. Act on liveness conditions engendered by policy
 - c. Accept, verify, and incorporate policy updates

5.4 Interactions

Two kinds of entities interact: agents and organizations. Organizations interact through the agents of which they are composed. Distinguishing between whether an interaction is agent-agent, agent-organization, or organization-organization can seem superfluous in a closed system, but becomes of paramount importance in an open system as an agent acting on behalf of an organization may have the authority to make commitments far beyond its local purview. For instance, a gateway agent might be authorized to grant access to an information repository or grant machine cycles for program execution; the agent’s organization would ultimately be the responsible party.

Performing Gaia analysis is a two-pass operation: The first pass captures necessary agent-agent interactions and the organizations involved; the second identifies how the organizations interact with one another and with individual agents. Note that were we actually executing the process we could not talk here about the interaction of organizations because we would not have identified any.

5.5 Organizations

In a closed system, the organizational elements are less noteworthy because the agents interact with one another directly and do not represent organizations. In our work with microgrids (see [Phillips]), we identified organizations called cells, globs, and co-ops. The general notion is that a *cell* is a set of sources and associated loads operated by a single agent; a *glob* is a group of self-interested cells that satisfy their own loads before any neighboring loads, and a *co-op* is a group of cells that may under some conditions satisfy neighboring loads before satisfying their own based on a shared policy. Note that this set of organizational types is useful for operating a closed system but also supports the kind of marketplace behavior needed to operate an open system.

5.6 Liveness and Safety Conditions

Liveness and safety conditions are the embodiment of system operating policy and define the desired behavior of the agents under various conditions. Liveness conditions are conditions that the agent system has the goal of bringing about; safety conditions are conditions that the agent system has the goal of maintaining. As [Zambonelli] puts it, the liveness properties of a system ensure that “something good happens” and the safety properties ensure that “nothing bad happens.” In electric power generation and transmission, the safety conditions are the normal conditions, e.g., appropriate frequencies, voltages, and temperatures; and liveness

conditions describe the desired system behavior in response to departures from normal conditions, e.g., alert or emergency system conditions caused by contingencies. An example of a safety condition is maintaining the alternating current (A/C) frequency; the concomitant liveness condition is to return the system to the proper frequency when it deviates.

The following is a set of liveness and safety conditions developed for a general distributed electric power management system. Conditions 6, 7, and 8 illustrate the inclusion of CFD concerns to the agents' behavioral structure.

1. Human interaction
 - a. When a person with the appropriate authority issues a command that the system is able to obey, the system should obey the command and report that it has done so.
 - b. When a person with the appropriate authority issues a command that the system cannot obey, the system should report that it could not obey the command and say why.
2. Source control
 - a. When total load is in excess of the maximum that the system can supply in its current configuration, transition to a new configuration, if any, that can supply adequate power. Check first for stored configurations designated as capable for equivalent loads.
 - b. If there are no appropriate stored configurations, search for some.
 - c. When choosing a configuration to supply power, prefer configurations that:
 - i. Generate equivalent power at lower cost;
 - ii. Differ less from the preceding configuration;
 - iii. Have a lower system-wide average fraction of rated power being carried by all lines;
 - iv. Supply a thermal load that occurs within the appropriate time interval
 - d. If it appears that total load will at some future time exceed the maximum that the system can supply in its current configuration, search for other configurations in which the projected load can be satisfied. Record each such configuration in conjunction with associated load information and other information needed to select among configurations. Denote the configuration as capable of satisfying its associated load.
3. Load control
 - a. Maintain service to all loads.
 - b. When load must be shed, shed noncritical loads before critical loads.
 - c. Prefer supplying critical loads to shutting down sources for maintenance.
 - d. Prefer shutting down sources for maintenance to supplying noncritical loads.
 - e. If it appears that projected load will soon be greater than the system can supply, determine the order in which to shed existing loads and which loads should be shed.
4. Maintenance scheduling
 - a. Take components offline as required by their maintenance schedules.
 - b. As a component nears 90% of its mean time before failure (MTBF), assign it high priority for being taken offline.
 - c. Take any component that exceeds 90% of its MTBF offline for maintenance (may be overridden by 3c).
5. Distribution path
 - a. When a distribution path fails, compute the power flow for the remaining network and determine whether any of the remaining lines will be forced to carry more power than their rated capacities.

- If a line is carrying more than its rated capacity, and there exists some other line not in service whose placement into service will allow a new load distribution where no lines are overloaded, place that line into service. If more than one such line exists, choose the line for which the system-wide average fraction of rated power being carried by all lines is lowest when the system is placed in the suggested configuration.
- b. When the steady-state power flow through any power system element exceeds its maximum steady-state rating, transition to a new configuration that will bring all system power-flows into specification. Such reconfiguration may involve a prioritized shedding of noncritical heat loads or those that can be served by other heat sources.
6. FACTS long-term control
 - a. When an event occurs that has been designated a trigger for recomputation of long-term setpoints, notify the other members of the CFD group that this event has occurred and, depending on policy for that event type, either:
 - i. prepare local data for inclusion in a system state object, or
 - ii. prepare local data for transmission to group members who need it.
 - b. When data needed for recomputation of power-flow allocation has been appropriately marshaled for transmission, transmit it to the group members who need it.
 - c. When a complete data set as needed for recomputation of power-flow allocation is received, execute the complete-data form of the power-flow allocation computation using the complete data set.
 - d. When computation of power-flow allocation should begin based on timing considerations, but the data needed for computation is incomplete, execute the incomplete-data form of the power-flow calculation using the data set as-is.
 - e. When computation of power-flow allocation is complete, adopt the relevant output values (i.e., those that apply to the branch assigned to the FACTS device performing the calculation) as long-term control targets.
 7. When power-flow allocation recomputation has been triggered but communication with some group CFDs is not possible, adopt a reduced-group stance, including reduced-group data requirements, communication profiles, and algorithmic forms.
 8. When no communication with other group CFDs is possible, adopt and implement the standalone long-term control values for the assigned branch.

Two tables that appear in [Habur], entitled “Steady-state applications of FACTS” and “Dynamic applications of FACTS,” show the corrective actions that a standalone FACTS device can take to avoid or correct various undesirable power system conditions. These tables constitute a list of liveness and safety conditions for standalone FACTS devices. Since a UPFC¹⁵ contains both a STATCOM and an SSSC, entries in the “Corrective Action” column are potential responses of a UPFC if any of {STATCOM, SSSC, UPFC} appears in the “FACTS device” column. This applies to both Tables.

¹⁵ The FACTS devices being studied at UMR are UPFCs.

5.7 Conclusions about the CFD Agent System

A distributed agent coalition managing an electric power system is the locus for power system operational and security policy and can enforce both policies. The primary benefits of an agent-oriented approach are:

1. Separation of concern – the agents’ role is to do what agents are good at, and the CFD systems’ role is to do what the CFD system is good at, so, in theory, problems of cross-coupling, interference, and redesign are avoided.
2. Ready provision of distributed services – An agent framework that incorporates the appropriate social elements—communication, negotiation, group formation—is required.
3. Relatively straightforward policy enforcement – translation from human-readable system policy to a set of liveness and safety conditions that the agent system can enforce is relatively easy, again given the appropriate agent framework.
4. Integration of operational and security policy – violations of operational policy can trigger an enhanced security stance (more-frequent requests for identification, performing encryption where none had been required in the absence of suspected penetration, requirement of group signatures, etc.); conversely an enforced and well-thought-out security policy can ensure robust operation with graceful degradation in the face of security attacks and penetration.

Based on the material in this report, prior consideration of agents used for managing electric power, and significant support from the literature, it appears that agents constitute an enhanced context for cooperative FACTS operations by providing a host of framework services, including communication processes, cooperation utilities, enforcement of security and operations policies, security operations, cooperative fault isolation, distributed algorithm execution, data marshalling utilities, and real-time cooperative planning to respond to unexpected contingencies. At least some of these capabilities have been demonstrated, but the cooperative agent-based system that reliably and regularly executes them all doesn’t exist. Our primary contention here is that if there were such a framework, distributed secure cooperative grid operations would be much easier to realize.

6 Consequences of Failure in an Electric Power System

The types of failure consequence in an electric power system are the same, regardless of the failure vector, the size of the impact, or whether the system involves agents and FACTS devices or not. This section provides a description of the consequences to power systems of failures, whether they are from natural causes or from malicious intent.

6.1 Consequence Metrics

Regional control areas (RTOs, ISOs, etc.) and utilities use fairly standard metrics to track the impacts of power outages. The three most common measures are the frequency of outages, the average duration of an outage, and the impact of the outage in terms of both loss of power and energy. The data are based both upon historical information and modeling. These data are often put into statistical form such as the likelihood that a particular outage will occur and how long it will take. In this way, depending on the amount of data a particular area has, there can be different types of data for particular regions, customers, and likelihoods of different severity of outages. This information is also used in the market to sell different levels of expected reliability to different sets of customers. Moreover, similar metrics can track other important metrics such as frequency and duration of overload conditions on lines.

- Frequency – There are different measures of frequencies for failures, such as the number of anticipated failures/year or the probability of a failure at any minute. The data can be aggregate for the entire system, or particularized to a location.
- Duration – The duration measures the duration of failures and as with frequency measures can be represented and aggregated in different ways.
- Size of Curtailment – The size of the curtailment can include both the probable power loss (MW) and energy loss (MWH) and, as with frequency and duration, can be represented and aggregated in different ways.

6.2 Public Consequences

This section describes different kinds of public consequences.

- Health and Safety – Safety impacts cover power failure events that affect public safety. Although people do die from the secondary effects of catastrophic power failures, and a few dozen people are killed each year by contact with energized lines, power outages do not directly cause large numbers of injuries and deaths. This is primarily due to backup systems in critical places—hospitals, e.g.—placed there based on knowledge that even short interruptions would be fatal to some and the experience that power does occasionally fail. However, extended outages tax emergency systems, which by-and-large were not designed to operate for long periods, and affect the safety of those living in electrically heated or cooled homes “in season.” In short, infrastructures are adapted to deal with short outages but usually do not have sufficient resources for interruptions longer than about 12 hours. Extended interruptions may have additional significant negative social effects such as vandalism and looting.
- Economic – Economic impacts are clearly involved in power outages. For example, many factories have critical multi-staged processes in which interruption of power causes loss of production in an intermediate step that results in significant economic losses even for a short duration outage. The August 14, 2003, Blackout Report (see [NERC]) estimated losses from the outage to be in the range of \$4 – \$10 billion. Though this is an extreme case, and estimates are difficult, it is clear that economic costs associated with power

outages are significant. Even under brownout conditions, where power is not lost but loads are serviced at lower levels than nominal, equipment can fail and cause economic losses.

- Political – A third type of consequence is the loss of public confidence and effects on the public image of a company that causes (or is perceived to have caused) a major interruption.

6.3 FACTS-Specific Consequences

The following relates the discussion above to specific considerations related to how CFDs operate in the context of a power system. The consequences of failures in FACTS device operations range from local outages and power interruptions to introducing stability problems into the system that could potentially lead to large-scale power interruptions. These consequences can result from either power system or adversarial vectors.

- Switch Level Control (Destabilization/Deformation of Power Wave)—As discussed previously the switch level control involves taking dynamic control commands and utilizing power electronics devices (IGBT switches, inverters, etc.) as well as other power devices (capacitors, transformers, etc.) to modify the real and reactive power flowing on the line. The execution of the power electronics is through hardware. Failure of the devices or modifying the inputs, either of the control commands or analog inputs from the power line could affect the operation of these controls. Modification of the controls could have two different affects—the power levels could be changed or distorted waveforms could be introduced into the power line. In the first case, load imbalances could occur due to shifts in power flow to other lines when the affected CFD line flow changes. In the second case, distorted waveforms could create stability problems by inducing transients, which can trigger relay controls to open power lines, or damage equipment of end users.
- Long Term and Dynamic Control (Damage/Tripping Caused by Incorrect Operation [Overloads] and Inter-Area Oscillations)—As with switch level control, affecting long-term and dynamic control settings would cause imbalances to occur that could shift power to other lines and cause overloads on the lines. Again, this could cause local outages and interruptions and lead to stability problems. Inter-area oscillations occur when machines (generators, loads, etc.) from one part of the power system oscillate against machines from another part of the system [Kundur]. This can result in transient ringing waveforms induced in the power system that through positive feedback can quickly cause major power stability problems. These oscillations can occur when separate CFDs react and respond to one another's settings, in effect creating a competition between CFDs that creates ringing and feedback. Two possible ways to manage the problem are to ensure that all the CFDs agree on settings before changes occur and only one CFD executes local settings changes at any one time, but simulation and testing is required to verify this.
- SCADA Interactions (Failure to get proper system data, incorrect FACTS settings transmitted)—In standalone FACTS devices, the main data between the SCADA system and the FACTS device are the FACTS settings and the readings of the FACTS device line. The SCADA operator may or may not be able to remotely control the FACTS device, depending upon how the device is configured and the operating procedures of the company. In a CFD environment, each CFD requires data for the status of the system it controls to execute the optimization algorithm used. If any of the SCADA data are compromised, if the magnitude of changes were great enough, it would affect the result of the optimization algorithm or cause the optimization to fail to complete, which could result in incorrect operations, overloads, and stability problems.

- FACTS Control Interactions (Failure of Cooperative Operations)—The optimization algorithm in a CFD environment is executed in a distributed fashion between the various CFDs used in the system. The CFD environment is also used to negotiate interactions between CFDs. If either of these functions failed or was disrupted, it could result in the same consequences—incorrect operations, overloads, etc. discussed above. However, if control of a CFD device were obtained, more direct adverse consequences would occur due to the direct control, whereas with other types of failures described, the consequences would be indirect since no direct control of devices is involved.

6.4 Power System Consequences

This section considers consequences that affect the power system. Power loss to a piece of power equipment is known as an outage; power loss to a customer is known as an interruption.

- System Responds Properly – If the system is sufficiently resilient and has sufficient backup capacity, the failure may take the particular equipment out of service (local outage), but the generators and transmission systems will absorb the transient changes and no customer power losses (interruptions) will occur.
- Local Interruptions Occur – The system may be sufficiently resilient, but due to the topography of the power system (e.g., there is only one radial feed to a particular load), there may be a local loss of power in an area (interruption); however, the rest of the system will not experience outages or stability problems.
- Stability Problems Occur – If a critical transmission element or generator fails, there may initially be only a local outage and/or interruption, but the failure can lead to voltage or frequency stability problems as well as overloading, which linger in the system and can trigger more widespread outages that can eventually result in a blackout or brownout.
- Regional Blackouts or Brownouts Occur – When an interruption occurs over a region, it is known as a blackout; when power is not interrupted but voltage is serviced considerably below nominal levels over a region, it is known as a brownout. There can be different sizes of regional interruptions, from areas of a particular city or regions of a state. Usually the term blackout or brownout is reserved for conditions of significant size, such as the August 10, 1996, West Coast and August 14, 2003, East Coast blackouts.

— This page intentionally left blank —

7 Contingency Analysis

Power companies and control areas design their systems to anticipate the types of failures that can occur. Today's power systems are expected to withstand some level of failure without major effects on the power system. There are limitations to what can be accomplished by design, as recognized by the NERC and embodied in their "N-1 criteria" requirement that a power system be able to provide power after the loss of any one major generator or transmission element.

Analysis tools used by power operators are discussed below. Analysis of the loss of one generator or transmission element is a first-order contingency analysis; analysis of the loss of two elements is a second-order contingency analysis. Usually, no more than second-order analysis is performed because of the dramatically increased effort required to go further. For example, for a 100-node system there would be 100 first-order analyses, 9900 second-order analyses, and nearly a million third-order contingency analyses for each system load profile. It is prohibitively costly and time consuming to go beyond second-order analysis or to evaluate many different load profiles.

7.1 Contingency Analysis Types

Contingency analyses can be subdivided into long-term and real-time analysis. Long-term analysis is used to design the system properly and create contingency plans to respond to specific critical contingencies that may occur. Real-time analysis consists of utilizing many of the same analysis tools to estimate the state of the system and possible contingencies that may occur for the specific system conditions that exist at the time of the analysis.

- Power Flow – Power flow studies examine load profiles during steady state conditions to ensure that under specified maximum conditions that system overloads will not occur. Power flow studies are also used for planning as well.
- Short Circuit – Short circuit studies examine the impact of failures on the system such as the resulting current and voltage levels at the time of the failure. These studies help with better relay coordination, and with determining plans to mitigate particular failures.
- Stability – Stability studies examine the short-term effects of failures on the overall power system. In essence, many of the analysis methods consist of obtaining a series of consecutive power flow snapshots over a short duration by allowing parameters to change as failures occur and determining how these changes propagate to other equipment. Each snapshot is a power flow program that factors in the effects from the previous power flow. Because of the complexity and computational requirements, the accuracy of these studies is limited to very short time durations, and only heavy load conditions with large impact equipment ("N-1") are examined.
- Voltage Drop – Voltage drop studies ensure that the system does not have inherent voltage drop problems in which normal expected load profiles cause voltages to be consistently low (or high in some cases).
- Relay Coordination – Relay coordination studies determine that relay values are set properly to insert and safely remove equipment from service in a coordinated fashion with other relays, to prevent improper relay settings from causing failures or compromising safety.
- State Estimation – State estimation consists of all the tools used to analyze the current state of power systems, using actual information to perform some of the above studies including power flow, short circuit and stability studies to determine how to respond to specific conditions that occur, that were not analyzed by long term studies.

7.2 Anatomy of Contingencies and Failures

When a power system contingency happens, several degrees of difficulty can arise. Usually, system protective measures prevent problems from extending beyond the local areas where they first occur. When large outages do occur, however, they are based on local problems feeding one another to cause regional instability that, in turn, trigger a major blackout. Figure 5 (following [Kundur]) diagrams these changes by tracking the basic transitions in various states that occur in power systems.

7.2.1 Anatomy of a major regional blackout:

local contingencies

lead to

local stability problems, outages, and interruptions

lead to

regional contingencies

lead to

additional local outages and interruptions

lead to

regional stability problems

lead to

regional cascading failures and interruptions

The August 14, 2003, East Coast blackout is a good illustration because a clear sequence has been documented: local difficulties in Ohio led to a blackout affecting the Northeastern and Upper Midwestern US and Eastern Canada. Our primary intent in describing this is to illustrate the kind of sequence that leads to a major power interruption.

- August 14 loads for the upper Midwest and East Coast, including large power transfers between regions, were high but not unanticipated for that time of year.
- At 1331 (EDT), a 600 MW generator tripped.
- At 1402, a 345 kV power line tripped. The resulting outages for these events were local. However, these events caused system voltages to sag slightly and inadequate system reserves were available to relieve the situation. This created an impending voltage stability problem in which increasing loads or other failures could trigger further interruptions.
- At 1505, 1532 and 1541, three additional 345 kV lines became overloaded and physically sagged into trees as they lengthened due to thermal expansion.
- From 1539 to 1608, sixteen 138kV lines in northern Ohio tripped out.
- From 1608 to 1638, 508 additional generating units throughout the Midwest, East Coast and Canada shut down.
- By 1700, tens of millions of people were without power; many would remain so for several days.

This brief synopsis illustrates how a few relatively insignificant failures can lead to significant regional power losses. For further details see [NERC].

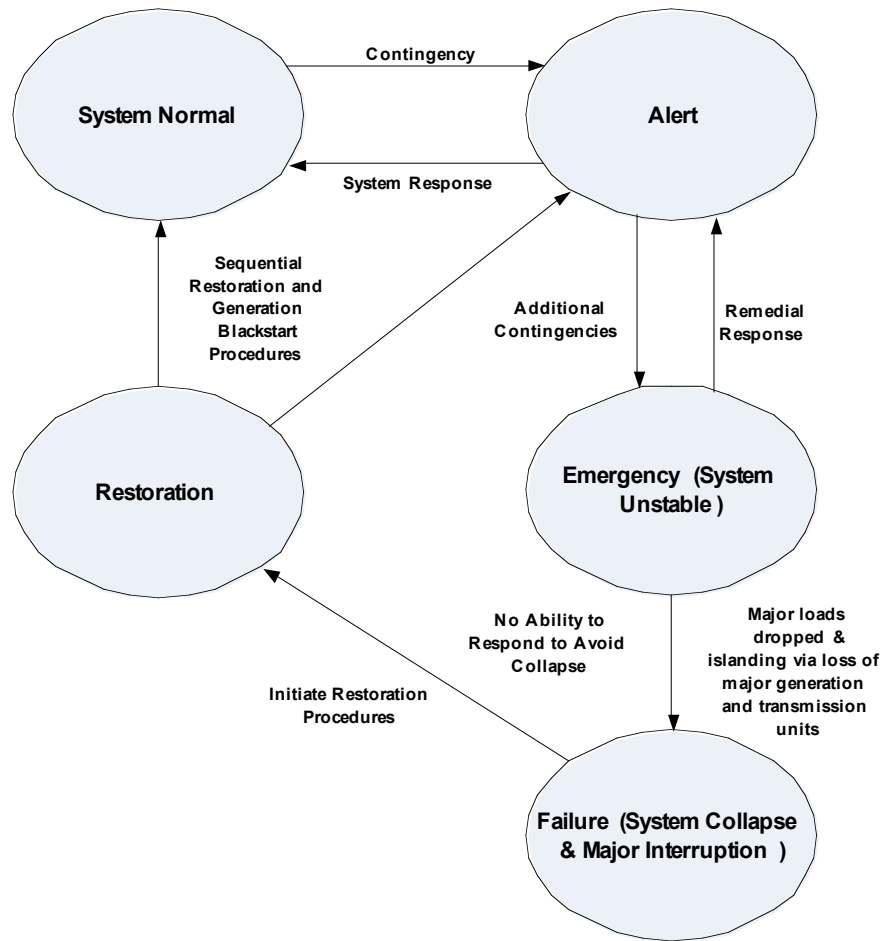


Figure 5. Illustration of Power System Transitions through various Contingencies

7.3 Natural Vectors for Power System Contingencies

The consequence associated with a power failure is due the type of contingency that occurs. The causes of specific contingencies are listed below. These causes are primarily from natural causes—lightning, equipment failures, etc.—and can also be caused by human error. These natural or human-error causes are referred to as *natural* vectors.

7.3.1 Normal Load Changes

Load changes both daily and seasonally in a regular but not entirely predictable manner. Under ordinary circumstances, load changes do not cause problems, due to well-defined processes utilizing automated systems relating generation to loads, market systems to bring new generation on line based on price and load, and the operational practices by which operators manage the loads in their jurisdiction. On the other hand, unforeseen situations can arise due to the complexity of the system and the great number of outside influences and individual customer and operator decisions. Occasionally these unforeseen circumstances can cause contingencies.

7.3.2 Generator Failures

Generators can fail due to a number of causes from mechanical failures, low voltages, and loss of synchronism with the system frequency. When a given generator goes out of service, the remaining system generators will increase output to pick up the load if there is enough system capacity available, but when the other generators are not co-located with the failed generator, the resulting power flows will be different even though the system is providing adequate generation. This may lead to overloaded lines in other areas. There may or may not be a local loss of power, depending on the topography of the system.

7.3.3 Transmission System Failures

Transmission system failures consist of all failures of transmission system equipment, from transmission lines, transformers, breakers, capacitor banks, etc., that cause equipment to go out of service. As with generator failures, the result will be shifting system power distribution and may result in local outages.

7.3.4 Control System Failures

Control system failures consist of all the failures associated with control of power system equipment. Since the control systems interact with and control the transmission and generation equipment, control failure can take equipment out of service or run it until it fails, mimicking or causing operational failure. The SCADA system is an example of a remote control system that can directly control large portions of the power system. Relay systems and local controls, such as those existing in standalone FACTS devices, are examples of local controls in which failures can cause equipment to fail to operate properly.

7.4 Adversarial Vectors for Power System Contingencies

A potential cause for all the contingencies listed in Section 7.3 is malicious intent, in which any of the natural vectors are purposely initiated by an adversary. The actual failures are identical to natural failures and the resulting consequences are the same. Malicious actions are referred to as *adversarial* vectors.

7.4.1 Denial of Cooperative Operation

A Denial-of-Service (DoS) attack is an attack on a computer system or network that causes a loss of service, typically the loss of network connectivity and services. A DoS is often accomplished by consuming the bandwidth of the target network or overloading the computational resources (resource exhaustion) of the target system.

While any connected system may be vulnerable to DoS attacks, systems working interactively, such as CFDs, are also vulnerable to a Denial of Cooperative Operation (DoCO) attack. The goal of such an attack would be to interrupt some or all FACTS-to-FACTS communication in an effort to partition or isolate the CFDs from one another. The affected devices would no longer receive the information from other CFDs needed to conduct cooperative long-term and dynamic control. An example of this kind of attack is to flood the information channels in a cooperative network with large numbers of inserted messages that are valid in form but contain nonsense or do not authenticate properly. Messages of this sort both consume network bandwidth and force the target system to waste computational resources determining that they are false.

Spoofing

Spoofing is the unauthorized use of legitimate Identification and Authentication (I&A) data, however it was obtained, to mimic a subject different from the attacker. Impersonating, masquerading, piggybacking, and mimicking are forms of spoofing.

A related attack known as a man-in-the-middle (MITM) attack occurs when an adversary reads, modifies, and re-inserts network messages between two parties without either party knowing that the link between them has been compromised.

If an attacker were able to successfully spoof a controller or another CFD, he could perform “trusted” communication sessions with a victim CFD. The victim CFD would believe it was communicating with the trusted source, but would actually be communicating with the attacker’s device. This could directly impact integrity or confidentiality.

7.4.2 Gaining Control over a CFD or Control Device

If an attacker were able to compromise a CFD or a control device, he could control the operations of the device and control or alter the messages communicated to other devices. This could directly impact availability, integrity, or confidentiality. This would probably be the most severe type of attack since the attacker would now directly control a CFD and could disable it, change settings, send out false readings, etc. Usually attacks of this type require defeating authentication methods and elevating privilege levels to be able access the equipment.

7.4.3 Desynchronization and Time-Based Attacks

A desynchronization attack is an attack against the temporal properties of systems that depend on synchronization for proper or accurate operation. The attacker uses some means to cause clocks to become desynchronized from one another, causing them to fail or operate improperly.

Cryptographic systems, once desynchronized, may take a substantial amount of time to resynchronize. Automated software and systems maintenance tools may make complex decisions based on slight time differences. The algorithms that will be used by CFDs are time dependent. The results are driven from feedback data from other CFDs. Data could be invalidated and long-term results skewed by incorrect time information.

7.4.4 Data Injection

In a data injection attack improper control or status information is injected into a system data stream. The general intent is to induce incorrect operation. If an adversary gains access to a communication channel and understands the communication protocol, he/she may be able to inject false but syntactically correct data packets. Data injection attacks *per se* do not necessarily involve spoofing, masquerading, using MITM, or gaining direct control of a device, but these methods can be used for data injection.

An example of a data-injection attack is the *replay* attack in which data is recorded and replayed at a later time on the same channel. If system communication protocols don’t include a mechanism to account for the time and/or order of the message, the system may be fooled into basing its data-dependent behavior on data from the past.

7.4.5 Malware Injection

Historically, malware (viruses, worms, trojans, etc.) has been emplaced in or targeted at well-known widely used applications (e-mail programs, browsers, and instant messengers), operating systems (Linux, Windows®, MacOS®), and devices (printers, routers, desktops).

There is nothing to prevent the development of specific malware targeted to specific systems if system information is available and methods to insert the malware are discovered. Properly designed malware can lead to partial or complete system cooption by an adversary.

7.4.6 Social Engineering

Social engineering is a blanket term for all methods used by adversaries to obtain information about a target system by exploiting natural human trust relationships. For example, an adversary may pretend that he/she has network problems in an attempt to obtain information about the type of network equipment and configurations used by a targeted system and use this information to gain access to protected system functions or to carry out one of the previously defined attacks.

8 FACTS Security Analysis

From a security perspective, a system of Cooperating FACTS Devices (CFDs) can be treated as any existing SCADA control system with one major distinction: CFDs incorporate elements of agent-based (or agent-like; see Section 5), peer-to-peer communication and control. The congruence of CFD systems to SCADA technology is advantageous because SCADA control system security is a well-researched area, and most of this research is applicable to CFDs. A key component of that research is the body of good SCADA practices, which offer many security practices that CFDs should employ.

Conversely, the agent-based, peer-to-peer aspects of CFD system operation are not as thoroughly thought out; applicable commercial security technology will be less common and relevant capabilities are more likely to be in the research phase.

Agent-based systems are a relatively young concept within computer science. As such, a number of competing methodologies, frameworks, and architectures for agent-based systems are under development. Although research into agent-based system security has produced significant relevant results, there is not yet a coherent, complete body of proven agent-based system security information. Therefore, a security plan for a system of CFDs may very well include unproven security elements based on ongoing research.

This analysis focuses on the security requirements, vulnerabilities, and implications when networks of FACTS devices coordinate their control actions. We have included comments and observations concerning the security of individual FACTS devices when appropriate based on the applicability of general SCADA system security principles.

8.1 Assumptions for this Analysis

Security analysis in this report is based on the following assumptions pertaining to the communication and control devices utilized in a CFD environment.

- CFDs communicate using IP (Internet Protocol).
- CFD computation occurs in the context of a commercial operating environment.

Protocols, operating systems, and communication mechanisms not consistent with these assumptions could easily be used. These alternatives are not specifically addressed in this report. Much of the report content, however, does not explicitly depend on these assumptions and therefore applies to such alternatives.

8.2 Security-Related Definitions

8.2.1 Security Element Definitions

8.2.1.1 Confidentiality

Confidentiality is the property of a body of information that it is available to only authorized entities and not otherwise disclosed. The confidentiality of a piece of information is enforced by ensuring that every access is properly authorized. Information cannot provide confidentiality; confidentiality must be enforced by some mechanism designed to provide it. A loss of confidential information may not directly affect a system but can cause major problems in other ways. The released information can damage a company or individual through public disclosure, provide advantage to a competitor or adversary, be used as a means for identity theft, or be used by an adversary as a precursor to attacks on integrity and availability.

8.2.1.2 Integrity

Integrity is the property of a body of information that it has not been altered by any unauthorized entity or mechanism. The integrity of a piece of information is enforced by ensuring that it has been changed by only authorized entities. An information system can be said to have integrity based on its ability to preserve the integrity of the information residing within it. A system's integrity depends on the correctness and reliability of its operating systems, the completeness and correctness of its hardware and software, the consistency of its data structures and processes, and the stored data itself. In a formal security model, integrity is interpreted to mean protection against unauthorized modification or destruction of information. Integrity attacks on an infrastructure control system are usually the most severe, because they can involve changes of system controls and data, which can cause the consequences described in Section 6 and the contingencies described in Section 7.

8.2.1.3 Availability

Availability is the property of a body of information that it can be acquired by an authorized entity as needed. Mechanisms that provide availability are normally required to meet timeliness and reliability requirements. Infrastructure control systems and their subsidiary information systems must generally meet information availability requirements no less stringent than those of the infrastructure itself.

8.2.1.4 Vulnerability

A *vulnerability* is a weakness that can lead to unauthorized activity. Exploits are related to vulnerabilities in that an *exploit* is the utilization of a particular existent vulnerability in carrying out an attack; in other words, a vulnerability is a weakness that can be exploited by an adversary. An *attack* is the series of steps taken by an attacker to achieve an unauthorized result using adversarial attack vectors (see Section 7.4), and may require multiple exploits to carry out. Vulnerabilities may reference any system element, including its information systems, security procedures, internal controls, or implementation [NISSG]. Note that vulnerabilities also exist due to normal operation, i.e., they may not be associated with any previously identified vector. Vulnerabilities of this sort are usually distinguished from those perceived during security considerations. [SAND2] provides a description of the kinds of vulnerabilities that exist in process control systems, including analysis of different categories for data, security administration, architecture, and platforms.

Table 4 lists the various categories of attacks that can be performed on CFD devices in a CFD environment mapping linking the previously discussed adversarial mechanisms to confidentiality, availability and integrity.

8.2.1.5 Threat

A *threat* is any circumstance or event with the potential to adversely impact an information system through denial of service and/or unauthorized access, destruction, disclosure, or modification of data [NISSG].

While the general definition of a threat covers any event, including natural disasters, this report focuses on man-made threats. In this context, it is valuable to consider the various groups of people who might wish to impact the system and what they seek to gain, because this may suggest the means by which they hope to achieve their ends and enable risk-based prioritization of defense activities (see the definition of risk; Section 8.2.1.6). In order to analyze a threat, one must take into account the threat source, including consideration of

Table 4. Categories of attacks on CFD devices

Attack Category	Attack Examples	Relative Degree of Difficulty
Confidentiality	Social Engineering	Low
	Scanning/Web searches	
Availability	Denial of Service (DoS)	Medium
Integrity	Spoofing and Masquerading	High
	Man in the Middle (MITM)	
	Data Injection	
	Desynchronization/Time Based Attacks	
	Replay	
	Insert Virus/Worm/Trojan	
	Gaining Control over a Device	

the specific motivations and capabilities of the instigator(s) and the timetable and steps required to carry out the attack. Table 5, from [NIST1], lists various threat sources, their motives, and possible threat actions they might take to plan and execute an attack.

8.2.1.6 Risk

Risk is the quantification of the possibility that a particular threat will adversely affect a target by exploiting a particular vulnerability [NISSG].

The following definition of risk is from [SAND1], which contains a guide for understanding and assessing risk from an overall system level perspective. The Risk equation is:

$$R = C \times T \times V$$

where:

R = Risk associated with an attack and/or system/asset failure

C = Consequence(s), the negative outcomes associated with degradation or failure of the system or asset(s). Consequences of an attack can be measured by loss of life, economic impact, loss of public confidence or other metrics

T = Threat, the probability or likelihood that a given attack scenario with the potential to disrupt systems or assets and cause undesirable consequences will occur. Threats are characterized by their means and likelihood of occurrence

V = Vulnerability, a weakness in the system or asset, or supporting systems or assets (e.g., security systems, etc.) to the threat (**T**) that would cause degradation or failure.

Consequences are the outcomes of the natural or adversarial vectors (vulnerabilities) that create contingencies that can lead to these failures. In order to quantify overall risks, each of these factors must be carefully considered. [NIST1] provides additional information on risk management.

Table 5. Human Threats

Threat Source	Motivation	Threat Actions
Hacker, Cracker	Challenge Ego Rebellion	<ul style="list-style-type: none"> . Hacking . Social engineering . System intrusion, break-ins . Unauthorized system access
Computer criminal	Destruction of information Illegal information disclosure Monetary gain Unauthorized data alteration	<ul style="list-style-type: none"> . Computer crime (e.g., cyber stalking) . Fraudulent act (e.g., replay, impersonation, interception) . Information bribery . Spoofing . System intrusion
Terrorist	Blackmail Destruction Exploitation Revenge	<ul style="list-style-type: none"> . Bomb/Terrorism . Information warfare . System attack (e.g., distributed denial of service) . System penetration . System tampering
Industrial espionage (companies, foreign governments, other government interests)	Competitive advantage Economic espionage	<ul style="list-style-type: none"> . Economic exploitation . Information theft . Intrusion on personal privacy . Social engineering . System penetration . Unauthorized system access (access to classified, proprietary, and/or technology-related information)
Insiders (poorly trained, disgruntled, malicious, negligent, dishonest, or terminated employees)	Curiosity Ego Intelligence Monetary gain Revenge Unintentional errors and omissions (e.g., data entry error, programming error)	<ul style="list-style-type: none"> . Assault on an employee . Blackmail . Browsing of proprietary information . Computer abuse . Fraud and theft . Information bribery . Input of falsified, corrupted data . Interception . Malicious code (e.g., virus, logic bomb, Trojan horse) . Sale of personal information . System bugs . System intrusion . System sabotage . Unauthorized system access

8.2.2 Network type definitions

8.2.2.1 Client/Server

In a client/server architecture, the server's role is to accept client connections while a client's role is to connect to servers when necessary to provide or acquire information (See Figure 6).

8.2.2.2 Peer-to-Peer (P2P)

A peer-to-peer network is a network in which every participant acts as both a client and a server and participants share information about known peers (See Figure 7). This type of network can provide a much more robust organization than the client/server model. With every peer being a server, the network is significantly less centralized, and no one server acts as a communications backbone. In a well-connected P2P network, every peer is capable of communicating with any other, whether directly or through one or more intermediary peers.

8.3 A System of Cooperating FACTS Devices as a Target of Evaluation

The Common Criteria is a standard for "specifying and evaluating the security features of computer products and systems [Abrams]." This reference has been selected specifically because it describes an "investigation of an innovative application of the *Common Criteria* (1999) in research and development, rather than acquisition."

Strictly speaking, a system of CFDs cannot be evaluated as a security target because the CFD system does not have a protective function and does not offer a protection profile. A CFD system fit for use in modern cyberspace, on the other hand, would necessarily have protective functionality and would permit the derivation of a protection profile. The Common Criteria, in particular as utilized in [Abrams], offers a standard, well-thought-out process for considering a system from a security point of view. In addition, Abrams provides "an approach to countermeasures characterization derived from the Common Criteria." A Common-Criteria-based approach to security, though involved, would provide a solid basis for communication with any government entities of interest and has a measure of international acceptance.

These comments about the Common Criteria and references to [Abrams] are included to illustrate the involved nature of a comprehensive security plan.

In addition, the National Institute of Standards and Technology (NIST) has produced a specific system protection profile for industrial control systems, [NIST2], that has been specified from the generic requirements in the Common Criteria. It includes an integrated set of security requirements for operating policies and procedures; technology-based system components, interfaces and interoperability; and physical security. Many of the rules defined in the security functional requirements section of [NIST2] could be applied to a CFD system to produce good security policies, procedures and implementation guidelines. Both [Abrams] and [NIST2] follow the framework briefly described in the remainder of this section to delineate requirements for security.

8.3.1.1 Target of Evaluation (TOE)

The Target of Evaluation is the system whose security properties are under evaluation. The TOE section of a Common Criteria evaluation describes the system and defines the scope of its operation.

8.3.1.2 TOE Security Functions (TSF)

The security functions of a Target of Evaluation are the set of all the hardware, software, and firmware of the TOE that must be relied upon for the correct enforcement of the TOE Security Policy (TSP) and definitions of the functional requirements to secure each of these.

8.3.1.3 TOE Security Policy (TSP)

The security policy of a Target of Evaluation is the set of rules that regulate how assets are managed, protected, and distributed within the Target of Evaluation, and methods to audit, measure, enforce or otherwise ensure that security is maintained.

8.4 FACTS Information Assets

This analysis focuses on the security requirements, vulnerabilities, and implications when networks of FACTS devices coordinate their control actions. Although we have included comments and observations concerning the security of individual FACTS devices, the novel aspects of the FACTS research being done at University of Missouri/Rolla concern cooperating FACTS devices. The primary vulnerability of cooperating entities is their communication in establishing common goals and plans and exchanging state data. Therefore, this analysis focuses on the information being exchanged by FACTS devices over a communication network. Table 6 below shows the types of information exchanged in a CFD environment and described below.

Table 6. Types of information exchange in a CFD Environment

Information Type	Description	Where Used
Data (Remote)	Data used to execute distributed max flow algorithm	System Status Input from SCADA System; Algorithm processing in each CFD; Data sent between CFDs to complete optimization algorithm
Data (Local)	CFD local monitoring and setpoint execution	Locally between CFD and Power line it controls
Control Settings	Human inputs to CFDs such as parameter values, policies, or direct commands	Remote or Local inputs to individual CFDs
CFD Control – Status	Current status of CFD, (e.g. in service, out of service, standby)	Internal to each CFD & traffic between CFDs
CFD Control – Setpoint Plan	Agreed upon plan to change setpoints in CFD environment	Internal to each CFD & traffic between CFDs
CFD Control – Setpoint Execution	Acknowledgment that setpoints have been executed in each CFD	Internal to each CFD & traffic between CFDs
Program Code	Distributed programming to compute optimization algorithm	Internal to each CFD
Identity and Authentication	Tokens and method to authenticate CFD as well as fixed ID info	Internal to each CFD & traffic between CFDs

Cooperative FACTS-based control and management of an electric power network requires these primary information elements:

1. Data concerning electrical phenomena
 - a. Local data about the branch assigned to a FACTS device coming directly into the device from dedicated sensors.
 - b. Remote data about parts of a network for which a FACTS device has no direct sensors coming into the device from a SCADA system or other information distribution. This data are received through a communications network.
2. Control settings from human operators. This information is typically received through a communications network.
 - a. Commands
 - b. Parameter values
 - c. Policy
3. Communication among FACTS devices needed to effect long-term control.

This item differs conceptually from 1b in that “remote data” refers to grid information exclusive of FACTS devices and not necessarily having anything to do with long-term control, whereas this item—“3. Communication among FACTS devices ...”—is concerned with information from other FACTS devices specifically for CFD system long-term control. This includes information pertaining to the states of various system processes and components: communication, computational, and security processes; cyber events such as attacks and network outages; the utility SCADA system, if any; and the FACTS devices themselves.
4. Executable computer program code
5. Identity information about each device. This information is unique to each FACTS device and helps both other FACTS devices and human operators identify and work with the device for authentication during cooperative interactions. This may include:
 - a. Network IDs, such as MAC addresses and IP addresses
 - b. Serial & model numbers
 - c. Firmware and software versions
 - d. Hardware identification, profiles, and capabilities
 - e. Credentials such as certificates, cryptographic keys, etc.

8.5 Inherited SCADA Security

A distributed FACTS control system is likely to utilize an existing SCADA system for inter-CFD communication, and, if so, will inherit security attributes from its SCADA roots. These include communications, control architectures, vulnerabilities, and more. However, due to the extensive body of research covering SCADA systems security, distributed FACTS control systems can also use a wealth of existing security “best practice.”

8.5.1 Physical Layer Communications

A distributed FACTS control system relies heavily on communications networks and architectures to enable long-term control of the power grid. Since each FACTS device individually makes changes on time intervals ranging from microseconds to seconds, timely

delivery of accurate, trustworthy data from peer devices and SCADA controllers are essential. Therefore, security of the FACTS communication network is vital to operational success of the distributed control system.

The physical layer of the network transmits information both among FACTS nodes and between SCADA control systems and the FACTS nodes. A variety of physical networks can support this capability at varying costs. However, each option carries known security protections and risks.

SCADA security guidelines suggest that control systems should be located on closed networks. However, we will discuss a variety of options in order to highlight the risks of physical layer choices.

8.5.1.1 Open Network (e.g., the Internet)

Open networks, like the Internet, offer any distributed information system a large level connectivity for a low price. Broadband connections are readily available and very cheap. However, direct connections to the Internet are notoriously insecure.

All traffic into and out of the connection is visible to third parties and is often logged by intermediate service providers. Because of the dynamic nature of routing protocols, the path of connectivity between any two nodes is not guaranteed. As a result, neither quality of service nor throughput is guaranteed. All nodes on the Internet are vulnerable to directed DoS or Distributed Denial of Service (DDoS) attacks. All nodes may also be adversely affected by the second-order effects of DoS/DDoS attacks on other systems on the Internet.

CFDs are not just a distributed information system, but also a control system for the power grid. Because of the critical importance of the power grid to the nation's infrastructure, CFDs should not communicate in an unprotected manner over the open Internet.

8.5.1.2 Closed Network

Closed private networks offer some advantages over the open Internet. Routers and other network equipment can be configured to prioritize network traffic, establish quality of service, and create static routes between nodes. A closed network is usually more secure than the Internet itself, but it is only as secure as the weakest link in the network. For example, a closed SCADA network connected to a local-access network corporate (LAN) is no more secure than the corporate LAN. Nodes in a closed network may still be compromised by accessing the hardware physically instead of remotely. Finally, closed networks can be substantially more expensive to deploy than using the open Internet.

Most power companies have already deployed or are deploying fiber-optic connections to their substations. These networks may be closed or, at a minimum, segregated from the Internet and corporate LANs. Provided that enough bandwidth capacity is available on these networks, they may be a reduced cost network option for cooperating FACTS deployments

8.5.1.3 Virtual Private Network (VPN)

Virtual Private Networks (VPNs) offer a middle ground between open and closed networks. A VPN uses open networks to send data between nodes, but secures the data using strong encryption to provide confidentiality, prevent network penetration by attackers, and generally reduce the effectiveness of network reconnaissance. However, VPNs are vulnerable to first- and second-order DoS/DDoS attacks and are only as secure as the weakest link that possesses the encryption key. The primary advantage of the VPN is the cost savings of using open networks rather than building and maintaining a closed network.

8.5.1.4 *RF Communication*

Currently, local and regional control centers use microwave communications equipment to control substation equipment. Microwave communications operate in the frequency range between 1 GHz and 300 GHz. This communication is primarily one-way, with substation equipment receiving commands from control centers. This simplifies the communication bus because only the control center can send information.

Communication among CFDs would require transmission capability for each device on the network. Furthermore, an arbitration algorithm (TDMA, CDMA, GPRS, etc.) would have to be deployed so that each CFD had a designated “space” for its messages. It is possible that the increased overhead of arbitration and the transmission needs of the CFDs would not fit into the current digital throughput constraints of microwave communications as used by power companies.

8.5.1.5 *Power Line Communication (PLC)*

Power line communication systems use electrical wiring to transmit information. They are already in use and power system operators are already familiar with their operation. These systems have various incarnations: home automation, automatic meter reading, control systems for switches and transformers, and broadband over power lines (BPL).

Home automation and automatic meter reading systems typically operate at the local distribution leg. They offer low bandwidth and low throughput communications in the 20–200 kHz frequency range.¹⁶

Long-haul communication systems are used to control and monitor switches, transformers, and other electrical equipment. These systems operate in the 30–300 kHz frequency range, and are increasingly used as backups or when microwave (RF) or fiber optic communications are unavailable. The data throughput may not be sufficient for a distributed, agent based system, especially when it employs peer-to-peer communication.

8.5.2 Control Architecture

8.5.2.1 *SCADA-based Client-Server (Centralized)*

The networking capability that allows CFDs to cooperate can also allow for direct communication between CFDs and SCADA Human-Machine Interface (HMI) workstations. Using this feature, FACTS devices could be fed information from a user, request information from other automated devices on the SCADA network, or be configured remotely. As with all remotely configurable SCADA devices, secure authentication and information certification are primary security concerns. Figure 6 shows the general client-server relationships. Note that the clients do not communicate directly with one another.

8.5.2.2 *Multi-Tier (Hierarchical)*

This is similar to the client-server case above, but instead of one-on-one communication between a user and each CFD, the user communicates with intermediary devices, each of which communicates with a group of CFDs on the user’s behalf. The goal of reducing the amount of direct interaction between the user and the CFDs is to increase efficiency, with the user configuring or sending data to whole groups of CFDs at once. Security concerns are the same as those for the client-server architecture. The major flaw in this architecture is that the intermediaries are points of communication failure for whole groups of CFDs.

¹⁶ Source: http://en.wikipedia.org/wiki/Power_line_communication




Figure 6. The Client-Server Architecture has a central information source, i.e., the server

8.6 Vulnerabilities

This section describes particular vulnerabilities associated with a CFD environment due to the cooperative peer-to-peer nature of the CFD environment. This supplements the material in Section 7.4 on adversarial vectors and vulnerabilities.

8.6.1 Operating System and Application Vulnerabilities

A CFD can be impacted by vulnerabilities in its host operating system. Operating system vulnerability concerns are the same as those for any network-connected device running that operating system. Programming errors or oversights in the design, implementation, or configuration of any application running on a CFD could introduce additional vulnerabilities.

8.6.2 Timing Errors/Induced Delays

The distributed FACTS control system depends greatly on the timely delivery of accurate information. Dynamic control makes control decisions on the order of microseconds, while long-term control operates on the order of minutes. Compared to human operators, both of these time intervals are very quick and enable the FACTS devices to respond quickly to changing conditions. Relative to computer processing cycles, these intervals are slow, thus enabling intelligent processing before decisions are made. However, relative to networks, dynamic control operates at the same speed, while long-term control is slower. Therefore, network performance, timing, and the results of network-based attack can impact the operation of dynamic control. Since long-term control uses dynamic control results as an input, these network issues may also have second-order effects on long-term control.

8.6.2.1 Network Time

CFDs depend on technology to keep electronic clocks and timestamps synchronized across the network. In addition to CFD operation, accurate time synchronization is vital to many security functions, such as secure authentication and activity logging. For network-based time synchronization to achieve secure operation, it must accomplish two objectives: first, the server must authenticate itself to the client; second, it must not be possible for an adversary to replay synchronization messages, thereby setting a local clock back by an arbitrary amount.

The most common time synchronization technology in use over IP networks is the Network Time Protocol (NTP). Research on distributed systems using unpoliced, standard implementations of NTP (See, for example, [Bishop]) has showed that insertion of a rogue time server can drift the clocks between systems by seconds, minutes, hours, or days. Another approach is to use time synchronization broadcasts, such as the NIST WWV broadcast or GPS signals. These signals can be spoofed but require the attacker choose an appropriate location. Should NTP exploitation techniques be discovered, the physical location of the adversary might be less constrained. A good practice is to have multiple sources for time synchronization information.

8.6.3 Code Revision/Replacement

FACTS devices incorporate complex computers that run binary instructions, whether stored on a hard disk, flash memory, or a programmable chip. An attacker that can knowledgeably modify those instructions can bypass any built-in local algorithms (e.g., dynamic control) and safeguards and take full control of the device. Then, by maliciously coordinating with other FACTS devices, the attacker can affect the long-term control algorithms being executed by the CFD group, thereby impacting a larger section of the power grid. The complexity of this attack is admittedly high, but so is the potential risk to the power grid.

8.6.4 Trust Manipulation

Each FACTS device acts as an individual node in a distributed information system that operates according to an implicit trust model. Trust is implicit in that proper operation of the long-term and dynamic control algorithms assumes trustworthy data from peer devices and authentic commands from any higher-level controllers. In a research setting, the trust model can be ignored while refining the functionality of the network. However, in a live critical deployment, vulnerabilities in the trust model can lead to outsider or insider attacks.

For example, a simple attack on a trust model is the MITM attack. In a cooperative FACTS deployment, neighboring devices A, B, and C might trust data from one another and use an IP table for one another's addresses. An attacker sees the frequent network traffic between these nodes and recognizes the trust relationship. The attacker floods device C with bogus network traffic, denying service to C and disabling C's ability to respond to A and B. The attacker then spoofs device C's IP address and sends malicious data to devices A and B. As a result, the long-term control algorithms from devices A and B start responding to the malicious data and begin sending power level adjustments to dynamic control. Device C receives the changes from A and B, and begins its own adjustments. However, C is unable to send out its current state. The attackers continue sending malicious data, and the negative feedback cycle continues until one or more of the power lines controlled by A, B, and C reaches its limit and overloads. The overload may cascade into the regional power grid.

8.7 Good Security Practices

8.7.1 Firewalls

Host-based firewalls should be installed on every networked FACTS device to help prevent potential attacks from within the device's local network. Networked devices should be divided into security zones based on security requirements. A network firewall should be set up at the entrance to any given security zone to create a combined front line for a group of devices requiring the same level of protection and to minimize the threat of attacks from outside the security zone.

8.7.2 Access Control Lists (ACLs)

Access Control Lists (ACLs) should be used to limit communication originating at and destined for FACTS devices to only the routes necessary for proper FACTS communication. FACTS devices should be allowed to communicate directly only with SCADA administrative computers and other FACTS devices. A sophisticated but growing concept is *network enclaves*, in which the data and functions of a system are grouped into enclaves and segmented from one another by the use of subnets, VLANs, VPNs, firewalls etc. Access Control Lists are used to disallow traffic without the proper credentials attempting to enter the enclave.

8.7.3 Intrusion Detection

Passive intrusion detection should be used on each FACTS device host and at the entrance to each FACTS security zone. These systems should be configured so that they do not alter the flow of traffic to and from the FACTS devices under any circumstances.

8.7.4 Authentication

Remote administration of FACTS devices should be designed to require login using exactly one unique username and strong password per user. All remote login accounts should be set to allow the user the minimum level of privilege necessary to do his/her job. The authentication system should incorporate access control protocols that allow users with different levels of authority to have different levels of access. The system should also allow the authority/access level of a user to be changed according to the user's assigned duties.

8.7.5 Logging

Logging should be done using a scalable logging subsystem designed to prevent flood attacks on the logging system itself. In addition to providing possible indicators that an attack is occurring, logging also provides an audit trail to clean up after attacks or system failures.

8.7.6 Tamper Prevention

Executable program code should be signed using certificates and/or multi-factor authentication. This should be combined with tamper-proof seals and other physical countermeasures that disable in-person changes to the device by unauthorized users (See Section 8.8.3, Multiparty Strategies).

8.7.7 Encryption

Encryption protects the confidentiality of data by encoding the data in such a way that only the intended recipient can decode it. The original, unencrypted data is referred to as *plaintext* or *cleartext*, and the encrypted data as *ciphertext*. Encryption can assist in providing data integrity and authentication, but not without additional measures such as error checking and secure key attribution. Encryption, and for that matter any function of any information security system, can be subverted by exploiting weaknesses in the protocols by which encryption keys are generated and transmitted, messages are queued for encryption, etc.

8.7.8 Integrity Checking

Integrity checking involves all the methods used to verify the integrity of data messages, such as CRCs, one-way hash functions, digital signatures etc. This can also include more esoteric methods such as biometrics and multi-factor authentication.

8.7.9 Configuration Management

Configuration management applies to both control system and network equipment. Configuration management allows one to trace configuration changes, which can be used to trace out the integrity of the system.

8.7.10 Non-Repudiation

Non-repudiation means, essentially, undeniability. In information systems, the term normally refers to the quality of a message by which an encryptor cannot deny having encrypted it and a decryptor cannot deny having decrypted it. Non-repudiation depends on authentication being in place, and can be used for accountability and identifying a message's source for fault isolation and information forensics.

8.7.11 Redundancy

Redundancy includes incorporating additional backup systems (preferably operated in “hot standby,” in which they immediately take over without interruption when a primary system fails) and backup paths for data to flow to anticipate failures. Redundancy can also be applied to security with the notion of “defense in depth” in which a system is designed such that the breach of any one layer of security cannot compromise the system.

8.8 Agent-Based Security

8.8.1 Control Architectures for Agent-Based Operation

8.8.1.1 Peer-to-Peer (Autonomous) Control

A CFD needs significant information about the power grid when making its decisions. Peer-to-peer communication, in which CFDs throughout the power grid inform one another of changing situations, allows this information to be communicated directly from one FACTS device to another. Figure 7 shows the notional peer-to-peer architecture: system components communicate with one another directly with no need for a central server.




Figure 7. The Peer-to-Peer Architecture has no central information source; devices can get data only from one another.

Peer-to-peer communication requires a notoriously high amount of bandwidth, so implementers must take care to ensure that the CFDs do not overwhelm the network. One way this might be accomplished would be to partition the peer-to-peer network into hierarchical communication zones based on proximity, minimizing unnecessary communication.

When agents operate independently in a peer-to-peer network, it is vital that they be able to trust the information on which their decisions are based. Therefore, a robust mechanism must be in place to ensure trusted communication.

8.8.1.2 Mixed-Mode Control

In this architecture, CFDs interact cooperatively as in the peer-to-peer architecture, but a user has the capability to send data or directives to some or all CFDs. This provides all the advantages of peer-to-peer communication, but with a degree of human supervision. Security concerns are a combination of those of the peer-to-peer architecture and of either the client-server architecture or the multi-tier architecture, depending on how user control is implemented.

8.8.2 Security-Oriented System Behavior

With respect to security, the response or reactive behavior of CFDs can be modeled at a high level by a state table. This generalization would apply to each expected information asset interaction between CFDs. Because the exact sensory data, deployment architecture, communication system, and level of peer verification are unknown and may vary by institutional rollout, a pre-determined global state table is neither feasible nor effective.

Table 7 enumerates the set of reporting behaviors and serves as an example of the analysis required for each cooperative exchange. An individual FACTS device that is part of the cooperative network expects a behavior from a peer or group of peers. That expectation could be the result of its own current state, a message from the peer devices, a routine process, an authoritative command, or even a joint calculation such as Max Flow. The individual FACTS device also perceives some behavior through its local sensors, remote sensors obtained from the SCADA system, peer messages, and/or data from other authoritative source such as human operators. Sometimes the expectations and perceptions match up; sometimes they do not. This kind of logically complete analysis enables the FACTS device to act appropriately—that is, according to policy—in all circumstances.

Table 7. Enumeration of Reporting Possibilities

CFD Behavior:	CFD reports:	Other sources report:	Analysis
O	O	O	All sources are reporting correctly
O	X	O	The CFD is reporting incorrectly
O	O	X	Other sources are reporting incorrectly
O	X	X	All sources are reporting incorrectly.

This table demonstrates the unpleasant notion that accurate reporting cannot be taken for granted in any of these cases. For example, the first row and the fourth row are indistinguishable to the observer (in both cases all sources agree), except the system is not performing as advertised in the fourth row. In order to decide how to respond, in each case the situation must be evaluated against the implemented definitions of trust and results of

decisions about which reports should be considered trustworthy. Even in cases without conflict, likely the most common case, it cannot be assured that what is reported is what occurred. The most important recommendations based on this table are the following:

1. Deploy the CFD environment with CFD algorithms that incorporate measures to detect internal conflict about an information element's truth value, and
2. When there is such a conflict, respond immediately, according to a specified policy, to stop any ongoing damage, prevent using corrupted data (if any), and take appropriate action with respect to any improper or malicious activity in progress.

The specific policy should indicate what action is to be taken in response to unexpected events that affect security. Possible responses include, but are not limited to:

- Record the unexpected behavior and continue standard operation
- Raise an alarm
- Send commands to shut down or restart the suspect device
- Shunt communications around the suspect device
- Stop responding to communications from the suspect device
- Ignore all communication from the suspect device
- Do not trust information coming from the suspect device
- Do not trust information from any device

These policies would have to be developed based on an analysis of potential attacks and failure modes that could occur. A risk analysis should also be performed to rank the possible events on a likelihood scale so security policy implementation can be designed to respond most effectively to the riskiest attacks and failure modes. A more comprehensive analysis would look at the various types of adversarial vectors as listed in Table 4 (and further discussed for peer-to-peer CFD environments) coupled with the types of information exchange in a CFD environment listed in Table 6. An analysis could then be done for each type of communication to compare the consequences and impacts of the adversarial vectors that could disrupt the CFD environment. With this analysis in hand, policies could be constructed to anticipate and provide remediation for the contingencies it contains.

Figure 8 supplements Figure 2 in illustrating the locations of vulnerabilities in a CFD environment. The sections below on multiparty and agent-based strategies provide further discussion of what constitute good security policies for a CFD environment.

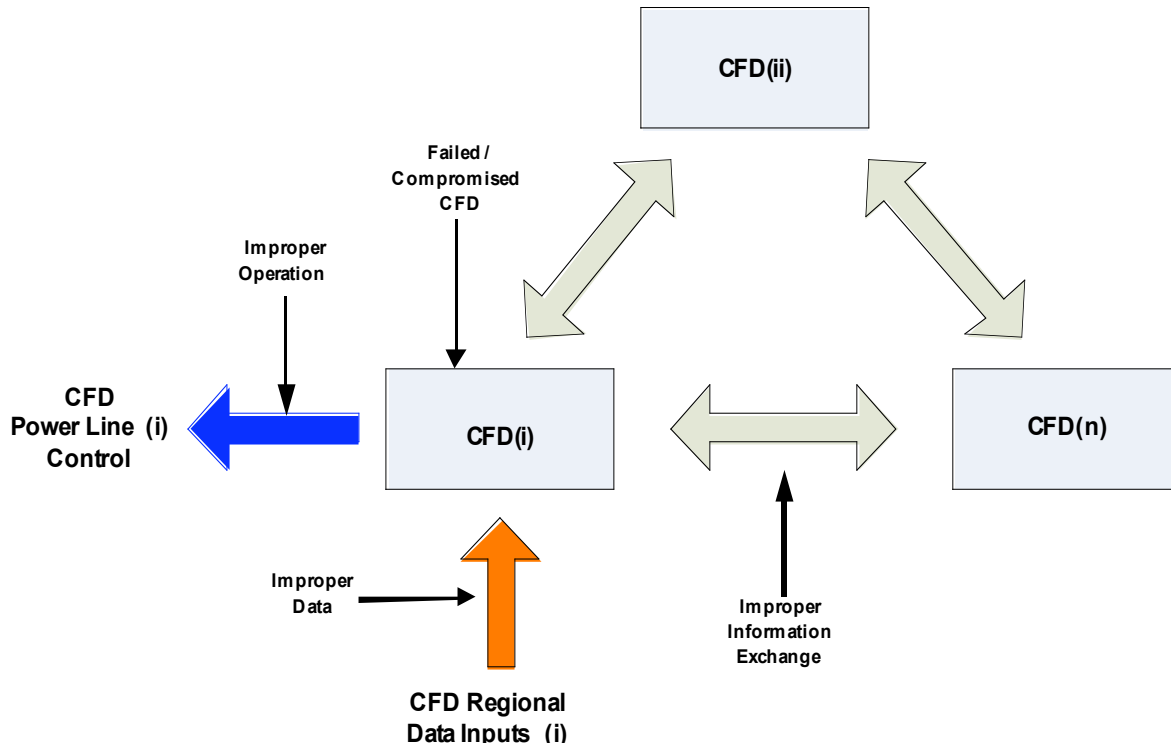


Figure 8. Location of Vulnerabilities in a CFD Environment

8.8.3 Multiparty Strategies

Failures that can be traced to human activity are often the result of a malicious act carried out in secret by a lone perpetrator or an error by a single individual. The likelihood of such failures would be reduced if more than one individual were required to participate. If any involved party did not want to proceed, he/she could stop the activity by withholding assent. Techniques that require participation by two or more individuals for successful execution are called *multiparty* techniques. Cryptographically implemented multiparty techniques are not explicitly agent-based, but are very complex, so that something like an artificial social system—such as an agent society—is helpful in implementing a system that executes them.

Multiparty techniques are used when the risk of negative consequence due to unauthorized unilateral execution is large. A familiar multiparty technique has been portrayed in many movies: In order to launch a nuclear missile, two keys must be inserted into two locks located several feet apart and turned simultaneously by two different individuals. This prevents launch by any single adversary. On the other hand, if either party does not act, for whatever reason, the missile cannot be launched, because even an authorized player can't do it alone.

For some multiparty cryptographic protocols (see, for example, [Goldwasser]), generating authorization requires at least n of m players, with n and m chosen to provide failure and threshold properties. For example, n less than m allows execution even when up to $m-n$ players abstain. The mechanism allows variation in n and m . The players would signify their willingness to authorize by interacting with a digital system (say, by entering their PINs). The complex mathematics requires a computer, which is also useful for keeping track of the state of the protocol. This again suggests (although it does not dictate) the use of agents.

For instance, suppose policymakers decide that at least three of five system administrators must be present when software is being installed. Multiparty authorization could be set up with, say, n being three and m being five, i.e., at least three administrators are needed to generate authorization to install. This technology-backed policy is intended to prevent any one or two individuals—in particular, malicious insiders—from installing software.

8.8.4 Agent-Based Mitigation Strategies

8.8.4.1 Fail-Safe Mode

Under various circumstances, an individual FACTS device may engage in a fail-safe mode of operation. Since FACTS devices are already operating in the field without cooperative behavior, an effective fail-safe option is to revert to standalone operation. This provides extra protection for the cooperative network because individual units can cease cooperation at any time to preserve their portion of the power grid.

Various circumstances that may require fail-safe operation include:

- Communications network failure
- Intrusion detection
- Safe mode command from an authorized controller
- Trust failure between cooperative peers

8.8.4.2 Manual Override

In the event a FACTS device requires maintenance, authorized personnel should be able to override cooperative behavior and force safe operation while upgrades and repairs are performed. This is to minimize risk both to the maintenance personnel as well as the electrical transmission network.

Any override mode carries the risk that it might be used maliciously by unauthorized attackers. This risk can be minimized by enforcing multifactor authorization; e.g., a CFD entering maintenance mode only when it receives a command from a trusted source *and* is presented with a local authorization credential such as a physical or electronic key.

8.9 Network Situational Awareness and Visualization

A powerful and important mitigation strategy against failure and attack of all types is to maintain a level of *situational awareness* at all times. This is accomplished by maintaining the state of a detailed model of the system of interest that reflects the current state of the real system. It is usually of interest that the system state reflected in the model not lag too far behind the actual system state. The amount of lag permitted depends on many factors.

Components are connected into larger networked systems that can exhibit collective behaviors not apparent from analysis of the individual parts. These emergent behaviors may be beneficial but can also create instabilities. The Sandia National Laboratories Center for Cyber Defenders has developed a prototype tool for the Analysis and Visualization of the Emergent Behavior of Distributed Intelligent Autonomous Systems (AnVEBIDAS) [Miller] that provides a generic agent framework and a number of visualization tools. The intent of this tool is not to automate the detection and classification of emergent behaviors, but rather to serve as a platform to visualize potential problems and the impact of system modifications. AnVEBIDAS has been applied to specific power grid scenarios (Figure 9 and Figure 10). We recommend that visualization technology be applied to the network communication system as well as the distributed power network.

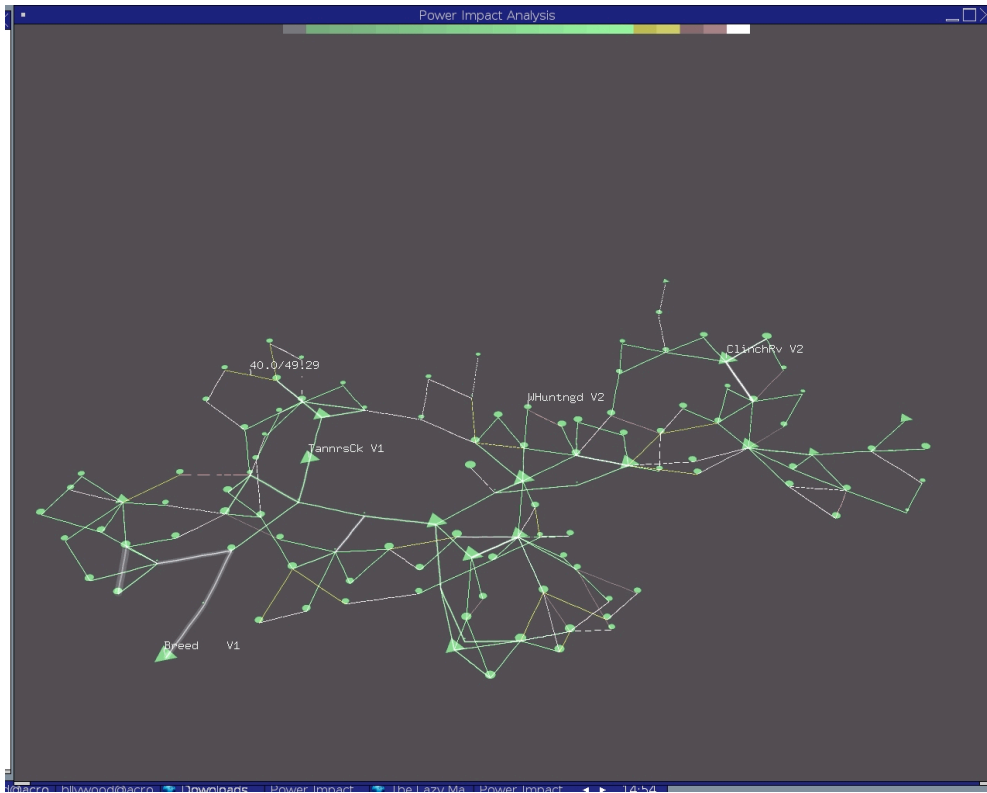


Figure 9. AnVEBIDAS screen showing the state of the IEEE 118-bus system

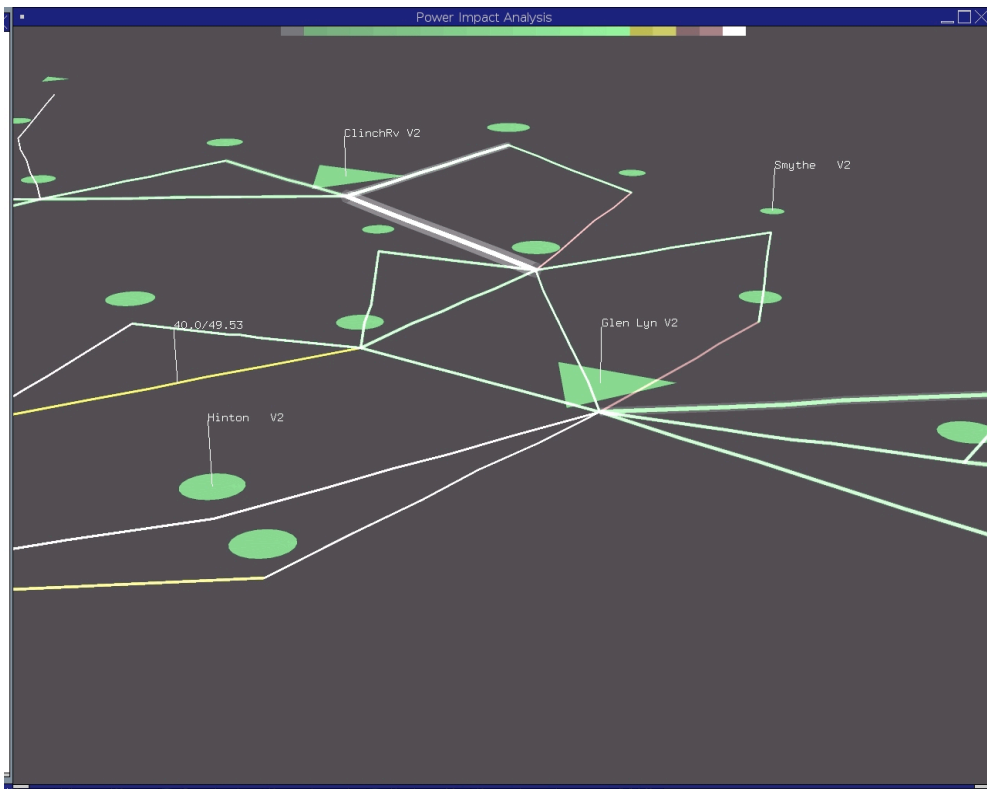


Figure 10. Detail of the AnVEBIDAS view of the IEEE 118-bus system

8.10 Security Documentation

8.10.1 Security Policy

8.10.1.1 Security Policy Elements

References [SAND3] and [SAND5] provide a framework for developing security policies for a control system. Figure 11, taken from [SAND5], illustrates the key elements of a comprehensive SCADA system security policy that includes elements for the different categories of a system that require policies for security.

[Fraser] provides an additional source for formulating cyber security policy. This document motivates the existence of a security policy and discusses its relationship to the security plan, the participants in its formulation, and key tradeoffs. The document recommends that security policy be implementable and enforceable using security tools and sanctions and that it clearly define the areas of responsibility for the users, administrators, and management.

According to [Fraser] the components of a good security policy include computer technology purchasing guidelines, a privacy policy, an access policy, an accountability policy, an authentication policy, an availability statement, an information technology system and network maintenance policy, a violations reporting policy, and supporting information that provides users, staff, and management with contact information for each type of policy violation and guidelines on how to handle outside queries about a security incident.

The authors recommend [Fraser] to anyone considering the acquisition or development of a security system; for novices it provides a straightforward and comprehensive guide and for experts a well-thought-out checklist and sanity check.

8.10.1.2 Guidelines for Writing Security Policy for CFDs

Every SCADA system must have a security policy in place. A good security policy is a code of standards and behavior with minimal reference to technology and no references to specific technologies; it should be able to survive major organizational overhauls with little or no revision. Installation of new technology should lead only to security policy revision in the extremely rare cases in which the new technology introduces to the system a new form of behavior that has yet to be codified in the policy document; in the case of CFDs, agent behavior may qualify (depending on its implementation).

Specific guidelines for how to apply the security policy in a practical manner should be addressed in the more specific security documents discussed in Section 8.10.2.

8.10.1.3 Agent-Based Security Policy

An agent-based security policy for CFDs is very similar to the SCADA security policy, but focused particularly on agent-level interaction, which is expected to be peer to peer. The included policies would represent a formalization of the state analysis performed using the technique identified in Section 7.5.3 in addition to any multi-party and mitigation strategies.

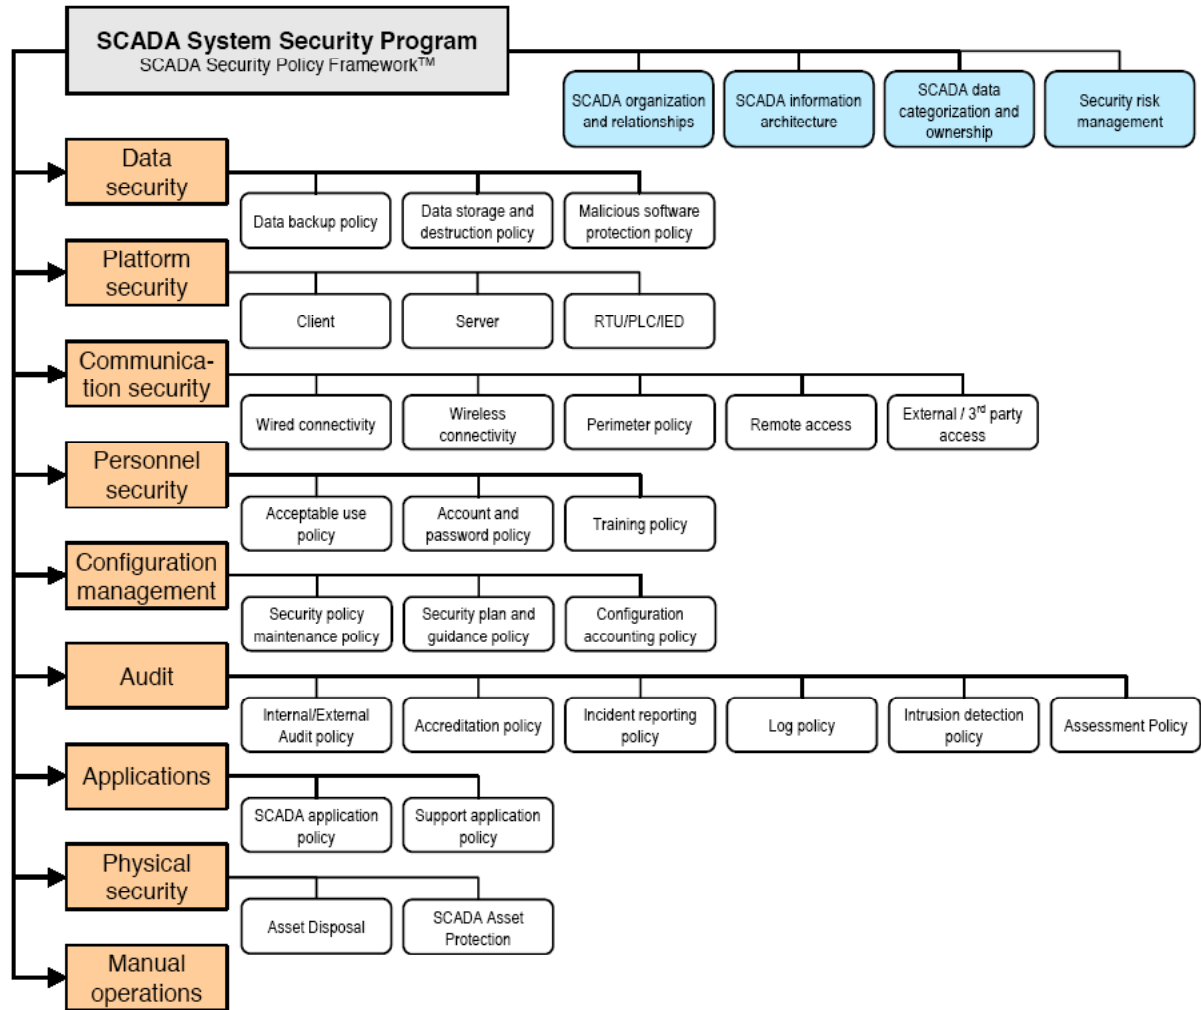


Figure 11. Sandia National Laboratories Security Policy Framework

8.10.2 Other Important Security Documentation

Like other components in the SCADA network, CFDs must be addressed in the network security documentation. While security policy is meant as an unspecific, relatively static guideline to secure operation for a system in general, there should be three other sets of security documents that are both more mutable and more detailed, and therefore certainly in need of amending to accommodate CFDs: the security plan, which applies the guidelines laid out in the security policy to the specific architecture of the network at hand; the implementation guides, which provide specific instructions for installing, configuring, and maintaining the specific system components; and the procedures document, which documents appropriate behavior for repeatable situations.

Because these three types of documentation must be custom-written for a given system, and because there are still many details missing with regard to CFD implementation, we can give very little guidance as to how these documents must be modified to include CFDs. As agent behavior is the only factor that differentiates CFDs from typical networked SCADA devices, most of the changes to the documentation will be those that are typical for adding any new

networked devices to the system; therefore, our discussion was confined to changes that may be brought about as a result of the agent-based nature of the system. Figure 12, taken from [SAND3], illustrates the relationships between the IT business network framework, site security policy, the security plan, and system implementation, and how they relate to various personnel managing different aspects of the control systems.

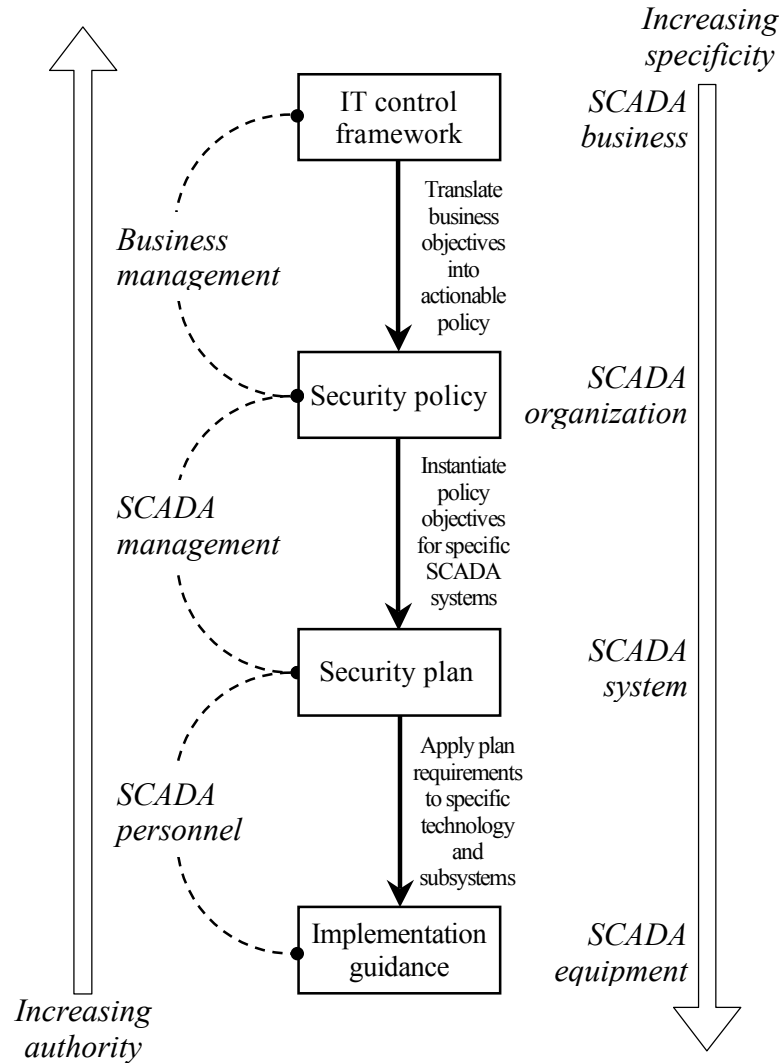


Figure 12. Relationships within the Security Administration Hierarchy

8.10.2.1 Security Plan

The security plan should apply the directives set forth in the policy to the types of components and behaviors known to be prevalent in the system. To modify the security plan to accommodate agent behavior, an organization must consider the potential for any installed agent's behavior to violate policy. This means studying the agent system's specifications and implementation and deciding which, if any, actions could result in a violation of policy.

An example of an appropriate security plan modification to accommodate agent behavior would be a stipulation specifying guidelines for logging of agent behavior in CFDs, thereby satisfying logging requirements that should be present in any security policy.

8.10.2.2 Implementation Guides

Implementation guides are normally very specific, requiring details such as make, model, version, and even firmware revision. The scope of an implementation guide should be the specific details of installation, configuration, maintenance, and removal. To account for agent behavior, implementation guides for CFDs should include details on appropriate agent configuration and instructions for installing, uninstalling, and upgrading agent software. Further, implementation guides for other network hardware and software may need to be modified to guarantee safe interaction with agent software.

8.10.2.3 Procedures

Procedures provide SCADA staff with a set of possible situations and appropriate responses to those situations. SCADA systems that include CFDs must be prepared with a set of procedures for agent misbehavior. Section 7.5 is a useful guide for potential situations of agent misbehavior, but guidelines for responding to these situations are very specific to the needs and priorities of a given organization and its SCADA system architecture. Like implementation guides, procedures must be extremely detailed and unambiguous, walking users through meticulous steps that were written specifically for the models and versions of hardware and software in the actual system. A different set of specific policies and procedures will also need to be implemented into the actual software code that executes the optimization algorithms to carry out setpoint changes and the cooperative agent interactions needed to carry out these procedures based on the security-oriented system behavior analysis discussed previously. So there will be a combination of procedures for monitoring, auditing and accessing particular CFDs in a CFD environment, and procedures implemented in the CFD software to respond to failures and potential attacks in real time.

9 Conclusions

This report considers the operational and security aspects of implementing cooperating FACTS devices (CFDs) utilized in an electric power system network (CFD environment). The first portion of the report addressed the operational aspects of a CFD environment and the second portion addressed the security considerations involved with a CFD environment.

The development of a CFD environment presents an opportunity to address the challenges posed by changes that power systems have undergone in recent years. Historically power systems have been physically and operationally isolated systems with few interconnections outside their own system. Power companies and regional organizations have been historically focused on reliability and cost issues. Security issues have become a new critically important issue for the following reasons:

- Power systems are becoming increasingly interdependent and interlinked
- Power systems need to manage changes brought by deregulation and open markets for generation as well as new regulations being developed by NERC and FERC following the August 14, 2003 blackout
- Power systems need to be able to adapt and respond to increasingly stressed systems with more limited opportunities for installing generation and transmission posed by costs and environmental and social concerns
- In addition, with advances in information technology the associated control systems associated with power systems have become increasingly automated and interlinked – legacy systems are increasingly being converted to COTS/internet based systems
- A wider group of people can access the cyber controls of these systems
- As a result of these changes, plus an explosion of hacker and terrorist threats on IT based systems, new vulnerabilities and risks now exist to CI's utilizing control systems
- Awareness, incorporation of new security practices, standards and technological developments must be implemented to mitigate these vulnerabilities

Due to these factors, the implementation and use of CFDs in a power system poses new opportunities for addressing these challenges.

Two approaches that appear to be particularly relevant are agent-based management and improved visualization. Both are responsive to modern power system concerns and could have made a difference in large contingencies on record. *Agent technology* provides a framework for cooperative distributed operation and *modern visualization technology* provides a more comprehensive and widely shared understanding of the state of the system. Together these enable better situation awareness, faster response time, and more appropriate response to contingencies. In addition, both technologies complement and enhance standalone and cooperative FACTS operations and operator interaction with the FACTS.

As understanding of CFD environments matures, future work should address many of the elements discussed in this paper, such as the unique operational conditions created by a CFD environment, agent interactions, and creating security policies and procedures at both a system level integrated with existing security policies and at an operational device level to respond to failures and possible attacks that can occur in a CFD environment. This report provides an extended framework in which further work along these lines can be done.

— This page intentionally left blank —

10 References

- [Abrams] Abrams, M.; *Applying the Common Criteria to DARPA Needs*; Mitre Technical Report 00W0000015; March, 2000.
- [Armbruster1] Armbruster, A., McMillin, B.; and Crow, M.; “Controlling Power Using FACTS Devices and the Maximum Flow Algorithm”; *Proc. 5th International Conference on Power Systems Operation and Planning, ICPSOP-2002*; Abuja, Nigeria; December 2002.
<http://web.umn.edu/~ff/Power/Papers/ETI-Nigeria.pdf>
- [Armbruster2] Armbruster, A.; Gosnell, M.; McMillin, B.; and Crow, M.; “The Maximum Flow Algorithm Applied to the Placement and Steady State Control of FACTS Devices”; *Proc. 2005 North American Power Symposium*; October 2005. <http://web.umn.edu/~ff/Power/Papers/Tpower04.pdf>
- [Armbruster3] Armbruster, A.; Gosnell, M.; McMillin, B.; and Crow, M.; Power Transmission Control Using Distributed Max-Flow; *Proc. 29th Annual International Computers Software and Applications Conference*; July 2005.
<http://web.umn.edu/~ff/Power/Papers/COMPSAC05.pdf>
- [Bellovin] Bellovin, S.; et al., Eds.; “Security Mechanisms for the Internet”; *Request For Comments 3631 (RFC3631)*; Network Working Group; Internet RFC Archives; December 2003. <http://www.faqs.org/rfcs/rfc3631.html>
- [Bishop] Bishop, M.; “A Security Analysis of the NTP Protocol”; Sixth Annual Computer Security Conference Proceedings; pp. 20-29, Dec. 1990.
<http://nob.cs.ucdavis.edu/~bishop/papers/1990-ntpsec/index.html>
- [Chaloupek] Chaloupek, J.; Tauritz, D.; McMillin, B.; and Crow, M.; “Evolutionary Optimization of Flexible A/C Transmission System Device Placement for Increasing Power Grid Reliability”; Submitted to the *2005 International Workshop on Frontiers in Evolutionary Algorithms*; July 2005.
<http://web.umn.edu/~ff/Power/Papers/FEA2005.pdf>
- [d’Inverno] d’Inverno, M.; Luck, M.; *Understanding Agent Systems*; Springer-Verlag, Chapters 7, 8 ; 2nd Edition, 2004.
- [Fraser] Fraser, B., Ed.; “Site Security Handbook”; *Request For Comments 2196 (RFC2196)*; Network Working Group; Internet RFC Archives; September 1997. <http://www.faqs.org/rfcs/rfc2196.html>
- [Goldwasser] Goldwasser, S.; “Multi party computations: past and present”; *Proc. Sixteenth Annual ACM Symposium on Principles of Distributed Computing*; Santa Barbara, California; United States ACM Press; 1997.
http://portal.acm.org/ft_gateway.cfm?id=259405&type=pdf&coll=GUIDE&dl=GUIDE&CFID=52724188&CFTOKEN=60410377
- [Habur] Habur, K., and O’Leary, D.; “FACTS: For Cost Effective and Reliable Transmission of Electrical Energy”; Siemens Power Transmission and Distribution Group, 2001.

- [Kundur] Kundur, P.; *Power System Stability and Control*; EPRI Power System Engineering Series; McGraw-Hill; 1994.
- [McMillin] McMillin, B., and Crow, M.; Fault Tolerance and Security for Power Transmission System Configuration with FACTS Devices; *Proc. 2000 North American Power Symposium*, pp. 5.1-5.9; October 2000.
<http://web.umn.edu/~ff/Power/Papers/ETIpaper.pdf>
- [Miller] Miller, S.; Service, T.; Atkins, W.; *AnVebidas: a Generic Agent Framework for the Analysis and Visualization of Emergent Behavior*; Sandia National Laboratories Center for Cyber Defenders project report; June 2005.
- [NERC] “Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations,” U.S.-Canada Power System Outage Task Force, April 5, 2004.
<http://www.nerc.com/~filez/blackout.html>
- [NIST1] “Risk Management Guide for Information Technology Systems”; NIST Special Publication SP800-30; 2001.
<http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>
- [NIST2] “System Protection Profile - Industrial Control Systems Version 1.0”; Prepared for NIST by Decision Analytics; 2004.
<http://www.isd.mel.nist.gov/projects/processcontrol/SPP-ICSv1.0.pdf>
- [NISSG] *National Information Systems Security (INFOSEC) Glossary*, NSTISSI No. 4009, September 2000.
- [Phillips] Phillips, L.R.; Link, H.E.; Smith, R.B.; and Weiland, L.A.; *Agent-Based Control of Distributed Infrastructure Resources*; Sandia Technical Report, to be published Fall 2005.
- [Rehtanz] Rehtanz, C.; *Autonomous Systems and Intelligent Agents in Power System control and Operation*; Springer-Verlag, 2003.
- [Ryan] Ryan, M.; Marksoe, S.; Cheng, Y.; Liu, F.; and McMillin, B.; “Structured Object-oriented Co-analysis/Co-design of Hardware/Software for the FACTS Power System”; *Proc. 29th Annual Computers, Software, and Applications Conference*, July 2005.
<http://web.umn.edu/~ff/Power/Papers/SandiaModelReport.doc>
- [SAND1] Baker B. et. al.; *A Scalable Systems Approach for Critical Infrastructure Security*, Sandia National Laboratories Report SAND2002-0877, April 2002. <http://www.sandia.gov/scada/documents/020877.pdf>
- [SAND2] Stamp J.; Dillinger J.; Young W.; Depoy J.; *Common Vulnerabilities in Critical Infrastructure Control Systems*, Sandia National Laboratories Report SAND2003-1772C, May 2003.
<http://www.sandia.gov/scada/documents/031172C.pdf>

- [SAND3] Stamp J.; Campbell P.; Depoy J.; Dillinger J.; Young W.; *Sustainable Security for Infrastructure SCADA*; Sandia National Laboratories Report SAND 2003-4670;
<http://www.sandia.gov/scada/documents/SustainableSecurity.pdf>
- [SAND4] Berg M.; Stamp J.; *A Reference Model for Control and Automation Systems in Electric Power*, Sandia National Laboratories Report SAND 2005-1000C
http://www.sandia.gov/scada/documents/sand_2005_1000C.pdf
- [SAND5] Kilman D.; Stamp J.; *Framework for SCADA Security Policy*, Sandia National Laboratories Report SAND 2005-1002C; 2005.
http://www.sandia.gov/scada/documents/sand_2005_1002C.pdf
- [Vlachogiannis] Vlachogiannis, J.G.; Hatziargyriou, N.D. and Lee, K.Y.; “Ant Colony System-Based Algorithm for Constrained Load Flow Problem”; IEEE Transactions on Power Systems; 20(3): 1241-1249; 2005.
<http://ieeexplore.ieee.org/iel5/59/32048/01490574.pdf?arnumber=1490574>
- [Zambonelli] Zambonelli, F.; Jennings, N. R.; and Wooldridge, M.; Developing Multiagent Systems: The Gaia Methodology; in *ACM Transactions on Software Engineering Methodology*, 12(3):317-370, July 2003.
<http://polaris.ing.unimo.it/Zambonelli/PDF/ACMTOSEMAccepted.pdf>

— This page intentionally left blank —

11 Definitions

- Branch:** Unless otherwise specified, a *branch* is an electric power line.
- Bus:** An electric power voltage bus, also called a node.
- CFD:** A Cooperating FACTS Device, that is, a FACTS device cooperating with other FACTS devices. A FACTS device is a CFD if and only if it is part of a group.
- Control system:** The collection of devices used to monitor and control the power transmission network. May include, but is not limited to, the SCADA system (Supervisory Control and Data Acquisition), the Energy Management System (EMS), Automatic Generation Control (AGC), Remedial Action Schemes (RAS), telemetry media and devices, sensors, actuators and data storage devices.
- FACTS** Flexible Alternating Current Transmission Systems
- Flow:** Unless otherwise specified, the term *flow* refers to the amount of real power, usually expressed in watts, carried by a branch. The real power sent through the branch is proportional to the current flow in the branch. Flow is controlled by a FACTS device.
- Group:** FACTS devices cooperating with one another are a *group*.
- IGBT:** Insulated Gate Bipolar Transistor
- Long-term setpoints:** *Setpoints* established by the cooperation of a group of CFDs
- Network:** Unless otherwise specified, the term *network* is used herein to mean an electric power transmission network, which encompasses the generation, transmission, distribution, and load devices (generators, lines, transformers, switches, capacitors, etc.) that generate, condition, transmit, and utilize electric power.
- Network state:** For the purposes of this analysis, the *network state* is a set of flows, one for each operational branch of the network. The *network state* also incorporates the underlying network topology (connectivity), branch states (impedances), equipment states (in or out of service), and voltage, current and phase angles for system branches and nodes.
- Node:** An electric power voltage bus.
- Policy:** A network's *policy* is the set of network conditions that the network's managing and controlling entities are to achieve, the actions they are allowed to take in order to do so, and the conditions under which they may take those actions.
- Setpoint:** The value of a branch's *setpoint* parameter is the appropriate amount of real and reactive power for that branch. A setpoint can be specified in several ways: (1) a power level not to be exceeded; (2) a power level from which branch power may not vary by more than some specified amount or percentage; (3) two values between which the power level is to be maintained; (4) voltage, impedance and phase angle.
- Sink:** An aggregate electric power load made up of individual residential, commercial and industrial loads.
- Source:** An electric power generating device. May refer to several such devices considered as a unit, such as a multi-generator oil or natural gas power plant.
- Vector:** As with a disease, a vector of a contingency is a causal mechanism that triggers the consequence or failure.

DISTRIBUTION

20 Dr. Bruce McMillin
322 Computer Science Building
University of Missouri, Rolla
Rolla, MO 65409

20 Dr. Mariesa Crow
233 Emerson Electric Co. Hall
Rolla, MO 65409-0040

40	MS 0455	Laurence R. Phillips, Dept. 5615
2	MS 9018	Central Technical Files, Dept. 8945-1
2	MS 0899	Technical Library, Dept. 4536



