Telecommunications In Complex Systems Failure: The Electricity Blackout of August 14, 2003

John C. Hoag, Ph.D.
McClure School of Communication Systems Management
Ohio University, Athens, Ohio, USA

Abstract

On August 14, 2003, residents of the Midwest and Eastern United States experienced the most significant electricity blackout in the nation's history. Not only were the magnitude and duration large, the level and complexity of automation involved arguably was unprecedented for public utilities. Given data recently made public, this paper attempts to comment on the analysis and conclusions on the primary investigatory bodies, identify the best practices in network monitoring and management that exist in the electric industry that could be applied to telecommunications, and to identify the best practices in network monitoring and management that exist in the telecommunications industry that could be applied to the electric industry. The primary differences between systems in these industries is the treatment of system state and its use in larger control scheme.

Introduction

Numerous conditions converged on August, 2003 to create a "perfect storm" in the electric transmission networks in the Great Lakes region of the U.S. Local failures in northeast Ohio "cascaded" into an event that interrupted power to over fifty million homes for up to one week, totaling 61,800 megawatts (MW) of electric power from 410 generation units – an economic loss of perhaps $10 billion in the US and $2.3 billion in Canada.[1]

This paper contains an overview of the blackout as well as different assessments of causality. Government and industry reports form the foundation of this section, when is then viewed through the lens of current cause-finding analysis. Following this general section are specific discussions of the realtime monitoring and control systems employed in the electric and telecommunication industries, respectively. This report closes with recommendations both for technology transfer for "best practices" between industries as well as for future research.

Blackout Background and Chronology

Electricity customers in the U.S. are served by a single retail distribution company and may choose to purchase energy from a variety of generation firms; this choice is important to industrial firms, which were the driving forces for this form of deregulation. This theoretical market is constrained by the physical network of transmission facilities connecting generators and loads. As such, bulk electricity is sold in bulk, subject both to the normal constraints (i.e., "congestion") of the network as well as day-to-day nuances in facilities maintenance. Unlike in telecommunications, no technology is available to switch or route bulk telecommunication; instead, power flow are governed essentially by the physics of Ohm's law and Kirchoff's current law. Power flows are additionally guided by factors such as relative frequency and phase angle, which is a characteristic of the type of load.

In the U.S., a traditional electric company would have been considered "vertically integrated," owning retail distribution, transmission, and generation assets. Deregulation or restructuring of the industry has often required structural separation or divestiture of generation but not yet the creation of independent "transco's" or transmission companies. The U.S. is thus separated into regions of transmission coordination that have responsibility for enabling contractual energy flows but have limited authority over generation or load; their work relies on realtime information systems.

Transmission coordinators universally appear to require that their members conform to a principle known as "N-1." Member firms must remain interconnected and operational in the loss of any single generation or transmission asset; moreover, in state N-2, each company must have processes to return to state N-1. These concepts are embodied as sets of scenarios and accompanying models that show each is feasible.

The blackout began in the region known as the East Central Area Reliability Coordinating Area (ECAR), whose investigators have painstaking developed a chronology of events.[2] These eleven events describe failures (i.e., "trips") of generation units and transmission facilities within and between companies. The ECAR task force also assigned causality to this event, to be discussed in the next section.

The events of August 14, 2003 can be summarized as follows: a warm day, system at 95% of peak load that is inductive/reactive due to air conditioners; some transmission lines in region offline but not correctly modeled. ECAR's chronology of events focuses on transmission line outages, mostly under extreme load, but (initially) under acceptable load. The initial event was the loss of a generator east of Cleveland at 13:31. The joint international task force separated events into the following four phases.
1. 12:15-14:14 Normal degradation
2. 14:14-15:59 Computer failures
3. 15:05-15:57 Local transmission failures
4. 15:39-16:08 Collapse of transmission network

Specific computer failures relevant to this paper include the following:
1. 12:15 Transco state estimator problems
2. 14:02 Transco mis-interprets event
3. 14:14 Wireco alarm system fails
4. 14:20 Wireco loses consoles
5. 14:41 Wireco server fails, restarts at 15:08
6. 14:54 Wireco backup server fails
7. 15:46 Wireco reboots servers
(Transco refers to the regional transmission coordinator; wireco refers to the retail distribution company)

Analysis of the Causes of the Blackout

Direct causative factors include failures of generation, transmission, information systems, and operations. In addition to the direct causative factors, latent and indirect factors contributing to this event include the pattern of bulk transactions created by deregulation of the generation sector of the U.S. electric industry, the reliability and usability of information systems, the inadequacy of system state computation models, as well as inadequate processes and procedures for firms employed in generation, transmission, and distribution of electricity. The distinction between these two categories of causative factors is significant both to understanding the chronology of the vent as well as implementing preventative measures to ensure reliable product delivery to the customer.

Woods[3] and Reason[4] have advanced the notions that both latent organizational factors and improper use of automation contribute to situations where errors occur. Structural and systemic issues are large contributors to the cascading of events on August 14, as well, in hindsight, as in gaps in the functionality of the decision support systems involved. Likewise, the remedies and overall preventative measures have a scope beyond a single firm and software system.

The Joint U.S.-Canada Power System Outage Task Force published its determination of the causes of the blackout's initiation. The Task Force's conclusions are summarized below:
1. "Inadequate System Understanding." The regional electric distribution firm at the center of the blackout had an insufficient understanding of its own transmission network, specifically at its operating margins. This firm is the collection of three operating subsidiaries that are interconnected by transmission facilities. Specifically, the interconnected system was not modeled under all relevant contingencies.

2. “Inadequate Situation Awareness.” The regional electric distribution firm did not recognize the condition of its system at the time of the blackout. It lacked necessary telemetry data from those external firms interconnected with it. Moreover, central monitoring system software and its servers were frozen or offline throughout this event.
3. “Vegetation Management.” As transmission voltage sags, current increases, causing lines to sag intro trees. Transmission line owners are
responsible for trimming trees. As transmission lines “tripped” or removed themselves from service, current at that time traveled over remaining circuits – often in excess of rated limits.
4. “Inadequate Realtime Diagnostic Tools.” Capabilities to monitor telemetry are normally augmented by mechanisms for control. The regional electric distribution firm apparently provided no single system view for its facilities and interconnection; moreover, it relied on strip chart recorders for some telemetry, a categorically non-integrated method.

The GE XA/21 Energy Management System installation at the retail distribution company failed to display any alarms after 14:14, a key contributor to the evolution of the blackout event. The company had planned to replace the system, but at the time was running a supported but out-of-date 1998 version (the vendor provided phone support during the event). The blackout revealed a latent bug in routine AEPR (Alarm and Event Processing Routine), which caused a ‘deadly embrace’ whereby two processes attempted to access the same data element simultaneously (each locking the other out, and willing to wait indefinitely). Vendor and user state that the deadly embrace scenario was “high unusual and unforeseeable.”[5]

The Electric Power Research Institute, a consortium of electric utility companies, provided a novel explanation for the start of the blackout, a method known as “fast voltage collapse.”[6] In the presence of heavy motor (e.g., air conditioner) load, supplier voltages sag and motors will stall. The duration of such a collapse is too brief to be modeled by current software, yet researchers consider its effects sufficient to trip generators in large quantities. In this case, the fast voltage collapse occurred in Ontario (not Ohio). The validity of this interpretation would undermine the many conclusions of the task forces cited in this paper.

The ECAR regional transmission coordinator listed the lack of sufficient intervention to relieve line overloads as the primary factor to this disturbance, also listing concern over transmission lines that tripped while within their emergency boundaries. ECAR did not significantly comment on information technology issues directly, subsuming their interest under the primary cause assignment.

Interestingly, the impetus of the dominant U.S. regulatory and policy body is to address issues of vegetation management.[7] The Federal Energy Regulatory Commission (FERC) has been committed to competition the generation sector and industry self-regulation in the transmission sector, while treating retail distribution as monopolistic and under state regulatory purview.

Power System Management

While the breadth of information systems in the electric industry is immense, the common characteristic of the applications discussed in this paper is the treatment of state. This is the single “best practice” of the electricity industry that should be considered by telecommunications network managers. The term Energy Management System (EMS) encompasses System Control and Data Acquisition (SCADA) and Automatic Generation Control (AGC).

The primary goal of state modeling is termed security or protection, which may be interpreted as fulfilling reliable service by avoiding outages by means of prediction. State models consider both voltage and phase angle for transmission and generation network elements. Of course, what is done based on this state information varies by type of firm: to shed load, to reject an order, to redispatch generation units, and so on. To model state is potentially a complicated task, yet it must be performed in near-realtime in order to meet protection and economic goals. Thus, some models will be recomputed every two minutes, while others including contingency scenarios would run on 30 minute intervals.

One vendor[8] outlined the operation of a state estimator as follows:
1. Import telemetry data (mostly analog)
1a. Inspect network topology
1b. Convert input data to correct units
1c. Quality check of converted data
2. Iterative estimation (of remaining elements)
3. Calculation of load flows
4. Reporting

Monticelli[9] has provided a useful framework for studying the use of state in power systems. He notes that the scope of any state model contains some territories that are observable and some that are not. The first step in state modeling is network topology processing – an active state that acquires sensor data, converts them to proper units, compares them to operating limits, and to track rates of change. Topology includes the aggregation of like elements into buses. The next step involves consideration of bad data, which extends NTP as follows.

Modern state estimation is performed through a more generalized topology processor that is based on linear algebra. Overdetermined linear systems for state estimation contain a variety of constraint equations based both on facility limitations and telemetry data. Some parameters are treated as random variables for estimation, The author notes that a variety of methods exist for solving entire and incremental models, with a common goal of efficient operation. He further notes that given the demands of deregulation (such as company mergers and acquisitions), the role of estimation in modeling is becoming more important.

Monticelli declines to describe state estimator models as realtime, instead using the term quasi-static. He notes that significant lag can occur between data collection in the field and arrival at the application program; he suggests even tagging samples with timestamps from the Global Positioning Satellite system.

To put this application suite into perspective, the retail distribution company at the center of the blackout is responding by replacing its entire EMS, procuring a projection system for system-wide status, and implementing a state estimator that solves contingency problems at automatic 30-minute intervals.

Telecommunication System Management

Developers of telecommunications equipment and operators of telecommunications networks have taken a systematic approach to ensuring reliable operation. The focus of this section is on upper protocol layer functionality rather than physical services such as SONET. Certainly, the physical and link layers provide, respectively, a failover/backup capability in addition to protocol techniques for reliable frame delivery.

More interesting and more transferable to the electricity sector are tools and techniques of Simple Network Management Protocol (SNMP). SNMP provides the means for monitoring direct and indirect status of nodes and links, and status includes both performance and security data. Devices that are compliant with SNMP possess a management information base (MIB) of ongoing information as well as the privilege of sending alarms to a host. Network management functions are interoperable among vendors because they have been developed openly by the Internet community and published accordingly.

Network management software not only receives alarms but typically correlates events for display and interpretation. Such hosts may receive alarms but may also poll devices to query (and reset) counters. Such software may need to be organized in a hierarchy, with other instances charged only with managing a subset of network assets. Although network management transactions are sent in-band, their latency can be managed (subject to the health of the device under investigation).

The notion of state in these systems is transparent, the presence of alarms and the other asset performance exceeding a limit. No control action automatically exists within the standard, but software implementations may launch trouble tickets or issue pages, for example.

The telecommunications industry has developed and promoted virtual services and service level agreements. While an enterprise network is capable of network management, the technology must be expanded to accommodate service level management. Lewis and Ray[10] begin with applications and infrastructure performance requirements and describe an environment where a "manager of managers" provides an additional layer of functionality but retains SNMP-based monitoring software to interact with field devices (and other managers).

Net management tools have the capability not only to receive security alarms, they also have the ability to interpret sets of alarm conditions as more complex attacks. This feature alone merits attention by the electric industry, given its interest in cyber issues.

Traffic pacing mechanisms are crudely implemented in network layer devices and are adequately implemented, however, in lower layers such as ATM, Frame Relay, and Switched Ethernet. Also, the suite of tools described in this section may be sufficient to ensure static levels of quality but lack the flexibility to be 'tuned' efficiently and bill customers accordingly. And most importantly, this suite of tools does not have a fully formed sense of state in the present tense, lacking the concern for future states and the mechanisms to avert trouble.

Conclusions

Given the stronger interest in telecommunications by the participants in this conference, this paper has critically examined the problem domain and information tools of the electricity industry. Shared by both industries are a desire for management systems to be timely, to display alarms usefully, to filter or interpret alarms, and to expand gracefully.

The electric industry and its information systems provide some good examples for developing telecommunications enterprises. The notion of system state as both a description and predictor of system behavior should be of great interest to telecommunications firms. Firms that interconnect are held to extensive requirement for quality of service as well as cooperation among control operator staffs. The coordination firm takes a superior role and has the privilege of regulating "traffic" of a member firm for the protection and security of the entire network. Speaking only of packet- and frame-based telecommunication networks, the ability to "shed load" is easily performed but the system state variables and constraint equations remain to be developed in the literature.

There are two issues to be resolved before transferring state-based concepts to telecommunications. The first is geography: the electricity model requires a regional coordinator that is not unlike a wholesale or transit Internet service provider (ISP), but we recognize that centralization in general impairs scalability. The second is consideration of future traffic: electricity has an active market for buying and selling future energy, and current telecommunications service level agreements are bilateral and static.

The U.S. Department of Energy's Grid 2030 project aims to borrow some of the best practices of the Internet and telecommunications to provide seamless reliability in a competitive setting.[11] A standardsbased reference architecture will place vendors on an equal basis and accelerate the development of plug-and-play instrumentation and software. The "GridWise" architecture will distribute intelligence to the field to improve reliability of telemetry.

The concerns of one EMS vendor in the electric industry could be shared by telecommunications managers: all of these systems are hard, costly, and time-consuming to implement. Moreover, their operation requires vigilance they potentially are not useful until fully developed.

In closing, in the absence of network-enabled telemetry devices in the field, the presence of telecommunications network management tools and techniques would not have alleviated this blackout. Assuming that such telemetry always runs out-of-band, on optical fibers rather than transmission lines, we can expect a redundant telemetry network to operate even while the transmission assets are offline. In the future, this infrastructure will enable power system management systems from the technological benefits the Internet hs provided.

Author

John C. Hoag is Assistant Professor in the McClure School of Communication Systems Management at Ohio University in Athens, Ohio, U.S.A. He holds a Ph.D. in Industrial and Systems Engineering from The Ohio State University, where he also studied in the National Regulatory Research Institute. The author would like to thank Dr. David Woods of OSU and Mr. Robert Burns of NRRI for their contributions to this paper.

Dr. Hoag may be reached at j.hoag@ieee.org and (740) 593-0077.
His website is:
http://www.csm.ohiou.edu/hoag

1. U.S.-Canada Power System Outage Task Force, "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations," final report, April, 2004.

2.East Central Area Reliability Council, "ECAR Investigation of August 14, 2003 Blackout by Major System Disturbance Analysis Task Force – Technical Report," February, 2004.

3.Woods, D.D., et al, Behind Human Error: Cognitive Systems, Computers, and Hindsight, Wright-Patterson Air Force Base, Crew Systems Ergonomics Information Center, 1994.

4.Reason, James, Human Error, Cambridge University Press, 1990.

5.Public Utilities Commission of Ohio, Case 03-2325, "FirstEnergy Corp.'s Response to the Entry of November 25, 2003," March 1, 2004.

6.EPRI PEAC Monitoring Team, "Initial Analysis of August 14th Northeast Blackout," Electric Power Research Institute, released August 14, 2003.

7.Robert E. Burns, private correspondence, June 10, 2004.

8.Sell, John (Elequant, Inc.), "The Future of Real-Time Energy Management Systems for Transmission System Operators," presentation to the San Francisco Chapter of the Institute of Electrical and Electric Engineers, Power Engineering Society, January 15, 2003.

9.Monticelli, A., Electric Power System State Estimation, Proceedings of the IEEE 88:2, February, 2000.

10.Lewis, Lundy, and Ray, Pradeep, "On the Migration from Enterprise Management to Integrated Service Level Management," IEEE Network, January/February 2002.

11.Meyer, David, "DOE Research Activities in the Office of Electric Transmission and Distribution," presented to the Edison Electric Institute, Charlotte, NC, April 19, 2004.