Symmetricom®

How Time Finally Caught Up with the Power Grid

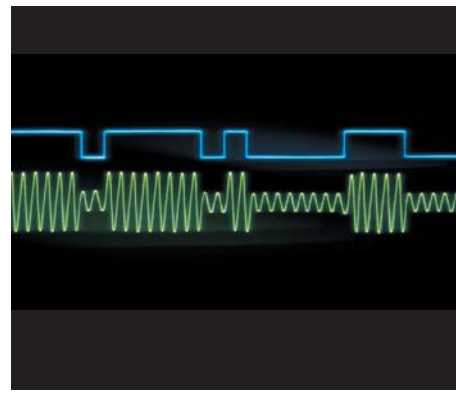

WHITE PAPER

How Time Finally Caught Up with the Power Grid

Besides electricity, regional power authorities must also learn to distribute synchronized time across the gird if they want to prevent blackouts like the one on August 14, 2003.

Depending on your point of view, massive power blackouts either happen in a few seconds or over several years. Take the outage of August 14, 2003. The actual collapse — a cascade of outages that affected 50 million people in eight states and the Province of Ontario — only took nine seconds. Yet, most experts agree that the conditions leading up to this event were years in the making — with consequences both foreseeable and inevitable. Those conditions included the inability of power operators to: 1) adequately assess the situation before the fact; 2) respond quickly enough once events started to unravel; 3) coordinate their actions in real time across the region; and 4) accurately timeline events after the fact and therefore pinpoint causes and prevent future outages.

The consistent theme across all of these factors is time — specifically synchronized time. Electric power supply and demand must always be in synch. Otherwise generators and transmission lines start to melt, burn up, or worse. Power itself is a synchronized event. Sixty times per second the direction of the electric current is reversed. If the cycle starts to speed up or slow down by just a little then generator turbines — designed to run in a finely tuned balance — begin to disintegrate.

The principles of electric generation are well known and have been since Thomas Edison built the first commercial generating plant (in Lower Manhattan in 1882). The principles of power distribution are also well known. What are less known are the principles of *time synchronized* distribution. Yet — as the massive power blackout proves — you can't distribute power reliably without reliable time synchronization between all the power stations. What is widespread synchronized time distribution? It is simply when every station and every piece of equipment in the power grid runs on the same clock, and the clock is correct.

The principle sounds simple — and it can be — provided you have the right "clock."

The 2003 Outage — What We Now Know

If nothing else, the event of August 2003 provides a clear benchmark with which to measure the reliability of the North American power grid following two key milestones in that industry. One of those was deregulation — an as-yet-incomplete process that has profoundly transformed the industry in the period since the mid 1990s. The other milestone was the Western power blackout of 1996 — which was actually two blackouts, the Western Blackout of July 2nd and the California Blackout of August 10th. As in 2003, what occurred were cascades in which a failure in one part of the grid brought down other parts within a few seconds. Before 1996, there hadn't been a cascade of similar scale since the Great Northeast Power Blackout of 1965. After almost 30 years of relative stability, there was good reason for the industry to think it had learned its lessons. The 1996 event was a wakeup call — one that the industry received precisely at the point when it was starting to reconstruct itself in the process of deregulation. If the system needed remedial work, this would have been a good time to do it. So what happened?

Months afterward the August 14th event, power industry regulators, grid operators, and the news media were still pulling together the facts. What was immediately obvious was that the cost and disruption to people's lives were enormous. Even though the exact tally may never be measured, common sense says it was in the billions. The California Blackout cost $1-3 billion and only affected a fraction of the people, industry, and real estate that the 2003 blackout affected. Whatever the statistics may finally show, it is probably the scenes on TV — of thousands of New Yorkers walking home across bridges — of five-star restaurants throwing out food — of families in Cleveland and Detroit lining up for bottled water — that best convey the blackout's impact.

If the effects of the crisis were obvious, the causes were anything but — especially to those actually running the affected grids at the time of the collapse. Nearly nine minutes after the blackout had swept from Michigan to Connecticut, system operators in the Midwest still had no clue how far the crisis had spread. "We want to know if anybody else is experiencing any problems," one official was quoted as saying.[1]

A month later, they and the rest of the world would have a much better view into what exactly transpired that day. Those details include the following:[2]

Once it began, the outage took only nine seconds to spread across the entire blackout area.[3]

The trigger event was the failure of a 345-KV power line near Cleveland that overheated, sagged into a tree and failed. [4]

At the time this line went down, eight other major transmission lines in the region had already failed and a number of plants in northern Ohio had been taken down that day, mostly for maintenance.

Even though these non-functioning assets were located in the same geographic area, oversight responsibility for them was split between two separate control centers — Midwest ISO in Carmel, Indiana, and PJM Interconnection in Valley Forge, Pennsylvania.

Twenty-two minutes prior to the blackout's start, controllers at these centers only knew of two of the failed lines — and only by sharing information by telephone.

Neither regional control center nor the power companies that operated this equipment had real time status information beyond the fact that the voltage in some lines was for some reason dangerously low.

An hour before the blackout Allegheny Energy, one of the utilities Midwest ISO had asked for more power, actually diverted power *away* from the ISO so that it could sell the power to other customers for more money.

When the ninth line failed, it created a 2,000 megawatt/10-second demand surge on the Michigan grid which in turn pulled on Ontario and Ontario pulled on New York, quickly collapsing all of them in a chain reaction.

The transmission company responsible for a large section of the Michigan power grid said that it did not even get a "courtesy call" from Ohio and was unaware of its neighbor's problems until two minutes before the collapse began — well past the "point of no return"[5]

Taken together, this is an astounding constellation of incidents — astounding because of the sheer improbability they would all happen at once, and astounding because so many of them were allowed to happen at all.

Why, for example, isn't real time line status information available to the controllers who need to see it? Why don't power companies publish their plant maintenance schedules online (just as they probably publish their conference room schedules online)? How could a supplier deny an emergency request from a neighbor in crisis just because it got a better price somewhere else? What's wrong with equipment that so much of it must be taken off line for maintenance at the same time?

The information vacuum in which controllers operate is blatantly obvious. The risks of operating in that vacuum are also obvious. The current system depends much more on luck than on informed management. A lot of things have to go wrong for the power to go out. The problem is — as the August 2003 event proved — a lot of things *can* go wrong — and that's increasingly likely to happen until there is corrective action.

The Bigger Picture

If the grid were indeed less reliable, that would imply that there is less margin for error and therefore a greater need for real time information about what's happening on the grid. But is the first premise true? Are power outages just bad luck — or is the grid inherently less reliable? The answer is yes, it is less reliable, a number of critics charge:

In 1999, the Electric Power Research Institute (EPRI) issued a report stating that the reliability of America's power gird is increasingly threatened while the technologies needed [to counter the threat] are delayed in commercial deployment due to a lack of financial incentives and R&D. ... This report finds that escalating demands for electricity, coupled with an outdated power delivery grid, pose a serious threat to the US economy.[6]

And there's this report in the *Wall Street Journal* following the August 2003 blackout:

> The main defects [of the North American Power grid] are transmission systems badly in need of improvement and a chaotic combination of regulated and deregulated markets. The unsteady regulatory situation, among other factors, inhibits the investment that is critically needed to improve transmission equipment. The nation is stalled in the middle of a massive but incomplete shift from old-fashioned state-regulated utilities — which generated their own power — to a system in which ownership of plants and transmission lines is broken up among a variety of players, and government oversight is fractured."[7]

The costs incurred from these grid defects go way beyond any single blackout — costs that are substantial and continuous. The Economic Policy Research Institute — a power industry think tank — estimates (based on its surveys) that annual economic losses from outages and other "power quality and reliability events" to be 1% of GDP, or $100 billion per year.[8] The defects result from a handful of fundamental root causes:

Under investment: Deregulation has forced companies to compete and therefore lower prices, leaving less money, and less incentive, to repair and replace assets, even as they are wearing out faster due to greater utilization and a load mix that is inherently more stressful on equipment.

Growing interdependence: Deregulation means power retailers shop around for supplies based on price — so buyers and sellers are much more geographically dispersed — and grid operators are forced to schedule deliveries across multiple grids in synch with each other.

A more dynamic environment: Prior to deregulation, generators might have typically been "spun up" once a day to meet the demand cycles of daytime operation or hot weather. After deregulation, generators are routinely taken on and off line on an almost minute-by-minute basis in response to the market. Market demand is much more variable in both timing and size (compared to, say, the sun coming up). This causes greater equipment wear, therefore a greater need for maintenance, as well as reduced margin for operator error.

A shift to digital loads: Independent of deregulation, today's grid load is becoming much more digital, as opposed to electromechanical in composition. In other words, the amount of electricity that goes into powering computers and other digital devices is increasing compared to the amount of power that goes into powering electric motors. That's significant because of the way computers and motors handle electricity. When a motor turns on or off, its load on the grid tends to wind up or down smoothly. When a computer shuts on or off, its load on the grid is almost instantaneous. As anyone who uses a computer knows, digital circuits are very unstable in the face of power interruptions or surges. But digital loads also make the grid less stable — which means that supply and demand across the gird must be much more highly synchronized to prevent voltage swings from getting out of control.

On this topic, *EPRI Journal* cites remarks by Carl Gelllings, EPRI vice president for power delivery and markets:

> Smart [i.e., digital] devices respond to very small changes in voltage by compensating in a way that makes the whole system operate in a non-linear manner. Such "non-linear" behavior causes a great deal of concern to those responsible for the electrical power delivery system because if they cannot predict how the system is going to operate, they cannot operate it as efficiently and reliably as desired. The system must therefore be operated with more margin and more allowance for safety. Another problem for the industry that arises as the use of the microprocessor increases is that equipment in general is much more sensitive to power quality.[9]

These problems have not come as news. Clearly, previous blackouts put us on notice. In fact, many experts say, the 2003 event occurred because (not despite) the fact that *some* of the remedies put in place after previous cascades actually did work as designed. Those include switches that automatically took generators off line in the face of load surges that would have otherwise destroyed them. What *also* should have worked were circuit breakers that were supposed to isolate sections of the grid *prior* to generator shutdown. If more sections could have been isolated prior to shutdown, even if only by milliseconds, then customers on those sections might have been spared the pain of a power outage.

For that to happen, however, a key requirement would have had to be met: Voltage sensors and circuit breakers would all have to be synchronized with each other across hundreds of square miles. In other words, there would be a distributed time on the grid, just like there is distributed power.

Given the conditions of the grid, the need for distributed time is actually growing much faster than the need for distributed power. Even the phrase "real time information" implies time is additional data that must be attached to whatever other information is being conveyed (i.e., that the information is current). And as the precision with which grid events must be synchronized increases, the term "real time" is defined in smaller and smaller increments. And as these units get smaller, there are more of them for operators to track in every hour and every day. So, yes, today's grid is less reliable and that does increase the amount of information controllers need. That's especially true for time: as timing becomes much more critical, as all information must now include timing information, and as the number of times that time itself is measured exponentially increases.

The Time Dimensions of Power

There are in fact two key roles time information plays in operating a power grid. One is to ensure supply, the other to ensure a balance between supply and demand. To ensure supply, grid operators monitor the cycle time of alternating electric current — nominally 60 cycles/second. If cycle time starts to drift higher or lower, that can indicate that voltage is too high or too low. A change in cycle time can also damage generators because turbines are correctly balanced to operate at that frequency. Not only does power have to be universally synchronized to 60 cycles/second — but those cycles have to be kept *in phase*. Otherwise, power flowing from one generator might actually work against power flowing from other generators. A 1998 Department of Energy report says:

> The vast, highly interconnected North American power system has been called the "greatest machine ever created." Generators separated by thousands of miles must rotate together with split-cycle synchronization, and the flow of power over thousands of transmission lines must be coordinated over large regions of the country.[10]

Not only do supplies have to be synchronized, but so do supply and demand. And they also have to be synchronized everywhere. For all practical purposes, electricity flowing through the grid cannot be stored. Once it is generated, electricity must flow *somewhere*. If there is not enough demand, it will cause voltage spikes; if there is too little it will cause voltage dips.

To maintain balance, generators and transmission lines must be turned on and off at the same instant that demand loads of equal size are placed on and taken off the grid. To coordinate supply and demand over wide areas, a large interconnected system of control centers has been created:

> The North American power grid comprises four major synchronous interconnections — western, eastern, Texas, and Quebec — which are further divided into ten Regional Reliability Councils (RRCs). Each RRC has primary responsibility for maintaining grid reliability in its region, which involves coordinating the activities of numerous control centers belonging to individual utilities, power pools, or — most recently — independent system operators (IS0s). Together, the RRCs compose the North American Electric Reliability Council (NERC), which sets overall reliability policies and standards.[11]

The job of these control centers, and the NERC itself, is to supervise a "superhighway for electricity commerce — carrying large amounts of power over long distances to ensure that customers have access to the least expensive sources of electricity throughout the year." The grid also acts as an emergency supply — allowing operators to import power from far away when local generating capacity can't meet demand.

Key to grid reliability, the Department of Energy report says, is synchronous operation:

> For a high voltage power network to remain stable, synchronism must be maintained. When this synchronism is disturbed by inevitable local events — such as a sudden loss of a major transmission line or a generator — power can begin to flow in an uncontrolled manner, causing automatic safety devices to "trip" and isolate parts of the system to prevent damage to equipment.[12]

Synchronous operation, the report goes on to say, is all about:

> ... real-time information needed for large highly interconnected transmission network, based on satellite communications and time-stamping [that enables] the system to operate closer to its limits.[13]

Synchronism includes:

> ... the ability to schedule power transfers on an hour-by-hour basis.[14]

When synchronism fails, so does the grid, stated NERC president, Michehl R. Gent in a CNN interview shortly after the August 2003 blackout:

So you have millions of people demanding electricity and you have dozens of power plants supplying that electricity. And it all has to be perfectly matched very second. And if something goes down and that match gets out of whack, then you can have a problem — like we saw last week.[15]

Distributed synchronous time also contributes to grid reliably *after* power failures occur. It gives operators a consistent timeline of what happened — information that can be used to prevent future outages. With thousands of events occurring in just a few seconds, precision timing of each event is crucial. Otherwise investigators won't be able to tell if one event is the cause of another event or the result. Following the August 2003 blackout, the NERC president was also quoted as saying: "Between 15 and 30 specialists at NERC are working to analyze data collected from every power station, utility and transfer hub that lost power." Events were calculated to several thousandths of a second and compared to time kept by an atomic clock.[16]

But the need to collect this data after the fact says two things about the system. First, it says that the information was not immediately available — and therefore not available in real time to controllers on the scene who might have used it to actually prevent a failure. Second, it says that whether or not timing information was available the grid nevertheless went out of synch — the grid did after all fail.

SCADA Requirements

So how do you synchronize a power grid? You start by saying which aspects of the grid you want to synchronize. Here's a sample taken from an article on measuring power system performance, published in *Electrical Construction and Maintenance* magazine:

> Your power quality monitoring system should be able to monitor a full range of performance characteristics and display useful information in a convenient form in real-time. The system must characterize steady state power quality irregularities, such as voltage regulation, unbalance, and harmonics, by using time trends and statistical disturbances. It must also characterize disturbances like transients, voltage sags, interruptions, and outages with detailed information about the events and provide trending.[17]

What's described above is part of a process known as SCADA, supervisory control and data acquisition. Currently power industry participants (regulators, power companies, transmission companies, high-volume users, etc.) are defining standards to perform SCADA across interdependent power grids. These standards will determine equipment and parameters to be monitored and controlled, also units of measure and communications protocols. The industry goal is to have a single platform from which any substation, generator, circuit breaker, or other critical device, can be visible to controllers without regard to geographical location or operating authority.

The most recent standard, and the one gaining the most industry traction, is DNP (Digital Networking Protocol) 3.0. One of many power companies to implement DNP 3.0 is Kansas City Power & Light. The reasons the company lists for its support of DNP 3.0 provides insight into how synchronized time fits into SCADA:

> KCP&L chose the DNP 3.0 protocol for the following reasons:
>
> DNP uses "report by exception" instead of traditional polling. Users may more effectively and economically manage communication.

DNP offers breadth and depth of protocol implementation. For example, report formats support time stamping for actual times in controls and relative to events.

KCP&L hopes to lower development costs of DNP applications by finding solutions for other utilities who will hare development costs and benefits

Future distribution automation projects with DNP will more likely be off-the-shelf instead of customized.

Standardizing on DNP eliminates the need for protocol converters.[18]

In other words, what power companies want is a plug-and-play SCADA technology that is available off the shelf to monitor and control events in real time. By implication, the search for this technology is a tacit admission that this is *not* where things stand now. Events (and the parameters that characterize them) are not related to time universally or consistently. Time is either not being measured or synchronism enforced. Time may not be part of whatever SCADA protocol is being used, or the sources of time are not accurate and therefore cannot be in agreement. In such an environment, the phrase "real time" is undefined because the grid does not know what the real time is. That means time stamps cannot be trusted, nor can the timelines and trend lines based on them. Since the grid must be, by definition, a synchronized environment that also means that grid failures become — literally — just a matter of time.

More Timely Line Fault Detection

Timelining events, scheduling power flows, and monitoring line frequencies — these are all applications that demand synchronized time. Another application is power line fault detection. Very much like radar, this is done by sending signals down a transmission line and waiting for an echo to return from a fault. The time it takes for the echo to be received must be precisely measured in order to determine exactly where the fault exists. Once again, precise timekeeping is required.

The search for a common SCADA platform also says something else — that whatever solutions are employed to solve timing problems today must cope with today's existing hodgepodge of interfaces and protocols. Regardless of which "clock" you set to turn on a power generator at a precise instant, and rotate at a precise speed, it had better have an output signal that the generator understands. A circuit breaker programmed to isolate a grid section within the first two milliseconds of a power surge had better receive a timing signal that is not only accurate but also completely unambiguous. But besides being backward compatible with existing connectivity and measurement standards, the timing solution would also need to be forward compatible with the new standard — be that DNP 3.0 or an alternative.

A Timing Solution Example

Increasingly, power companies and system operators are responding to these timekeeping requirements. They are installing time and frequency systems throughout their infrastructures that meet six criteria:

The timing system is *accurate* to within one millisecond. One millisecond is chosen because this is the resolution of the modern Sequence of Events Recorders.

The timing system is *absolute* rather than relative. This means that the time standard is related to a Universal Coordinated Time (UTC) source such as those maintained by the National Institute of Standards and Technology (NIST) and the U.S. Naval Observatory (USNO).

The system can be *operated unattended*. After installation, workers are not required to recalibrate the clock (in fact, they never need to calibrate the clock in the first place).

The system is *reliable*. The most accurate system does little good if it is not working, especially during a disturbance.

The system uses *standard hardware* for easier and less costly implementation and management.

The *cost* per installation is reasonable.

An example is the XLi Time and Frequency System from Symmetricom. The unit produces timing signals that can drive virtually any component in any power grid to within 30 nanoseconds of UTC. For power frequency applications, the unit offers an accuracy of 1×10^{-12} seconds — far more than enough to maintain 60-cycle/second phase synchronism. UTC is received from GPS satellites — so every device connected to an XLi, regardless of location, is running on the same master clock. That means there is in fact a true definition of "real time" across the grid.

The XLi supports both types of application timing requirements — i.e., to synchronize events so they happen on time or to record when they in fact do happen (to support timeline and trend line analysis). A wide range of option cards ensures compatibility with nearly every possible output/input requirement for time and frequency applications — just by combining up to ten options, oscillator upgrades and two power supplies per unit. Configuration recognition software automatically detects the unit's setup, without modifications to the operating system, providing plug-and-play configuration capability for current and future application needs.

Conclusion

The controllers who sit at the command consoles of today's regional grid authorities have jobs very much like FAA flight controllers. Just like airplanes, electricity must go *somewhere*, and also must go there *now* — except that electricity travels at the speed of light. Thousands of interdependent events *will* happen with split second timing whether planned or not. Generators, transmission lines, circuit breakers, power substations, transformers and a long list of other equipment are scattered across thousands of miles. Each responds to the beat of whatever clock it happens to be connected to. If those clocks are not synchronized, then control is left to chance. Sooner or later — as we saw on August 14, 2003 — those odds catch up to the grid. Solutions to prevent this from ever happening again are now readily available and easy to implement.

[1] "Overseers Missed Big Picture as Failures Led to Blackout," Eric Lipton, Richard Perez-Pena and Matthew L. Wald; *New York Times*; 9/13/03; Section A, P. 1.

[2] *New York Times*, 9/13/03, Section A, P. 1.

[3] "Blackout data will take weeds to analyze," Associated Press, 8/27/03.

[4] "Northeast blackout has now been traced to failure in Ohio," *Electrical Construction & Maintenance* website, 8/20/03.

[5] "Poor cooperation partially blamed for blackout," CNN website, 9/3/03.

[6] "How reliable is our electricity?" *Electrical Construction & Maintenance*, 12/1/99.

[7] "Overloaded Circuits: Outage Signals Major Weakness in U.S. Power Grid," Rebecca Smith; *The Wall Street Journal*, 8/18/03.

[8] "Estimating the Costs of Power Disturbances," *EPRI Journal Online*, 7/15/03.

[9] "Smart Power Delivery — A Vision for the Future," *EPRI Journal Online*, 6/9/03.

[10] "Technical Issues in Transmission System Reliability: A Position Paper of the Electric System Reliability Task Force," Secretary of Energy Advisory Board, 5/12/98, p.1.

[11] Secretary of Energy Advisory Task Force Report, p. 1.

[12] Secretary of Energy Advisory Task Force Report, p. 2.

[13] Secretary of Energy Advisory Task Force Report, p. 12.

[14] Secretary of Energy Advisory Task Force Report, p. 12.

[15] "Expert: Why wasn't blackout isolated?" CNN website, 8/18/03.

[16] "Blackout data will take weeks to analyze," Associated Press, 8/27/03.

[17] "Measuring Power System Performance for High-Reliability Applicatons," Mark McGranagham and Ross Ingall; *Electrical Construction & Maintenance*, 3/1/02.

[18] "KCP&L Enables DNP by Finding Missing Communication Link," Carl R. Goeckeler,"

SYMMETRICOM, INC.
2300 Orchard Parkway
San Jose, California
95131-1017
tel: 408.433.0910
fax: 408.428.7896
info@symmetricom.com
www.symmetricom.com