Design & Implementation of an Industrial Facility Islanding and Load Shed System

By

Michael Schiefen
BASF
Freeport, TX

Mark Adamiak, Bernard Cable
General Electric
Wayne, PA; Schenectady, NY

Abstract

The availability of electricity supply in all areas of life has become a critical need. In the industrial arena, costs due to the loss of electrical power and resulting process interruption can usually be quantified into hard dollars. Many larger industrials, however, contain significant co-generation capability that produce both electricity and waste heat for steam production utilized in the plant processes. At any instant in time, the available generation may be able to supply the needs of the plant if islanded; at other times and conditions, a sizable portion of the plant load must be shed in order for the load to match the available generation. In all cases, the goal is to maintain as much of the plant process as possible during an external disturbance or island condition in order to minimize environmental impact, production loss, and potential equipment damage.

This paper presents the design and implementation details of a control system that detects an island condition in an industrial facility, creates the island if needed, and executes a multi-tier load shed based on the load-generation balance that existed prior to the creation of the island. System operation in island mode and performance metrics obtained from dynamic tests are presented.

Business Case

As with all major industrial installations, reliability was a major consideration in the plant's design. With few exceptions, most components of the electrical system were designed to withstand one failure (N-1 contingency). Dual transmission lines from the Utility Grid were backed up by two 37.5 MW generators to provide power to the plant. This design allowed for the loss of any one source without impact to normal operation. All sources feed a 34.5 kV distribution bus in the main incoming substation. The distribution bus then feeds several of the plant's double-ended substations. Even at the utilization level most motors have installed spares. With this design, it was believed that the plants power system was secure and was capable of providing continuity of service under most failure scenarios.

Commercial operation began in December 2001. In May 2002 the security of the power system design was tested during its first major event. An insulator flashover in the plant's 230 kV substation ring resulted in a complete separation from the local utility. The site's cogeneration units were able to successfully operate in "Island Mode" for a period of approximately 12 hours. Eventually, the substation issues were corrected, and the Site power system was resynchronized to the utility grid without event or loss of power to the site.

One week later, the system experienced its second event due to an electrical equipment failure - again initiating "Island Mode" operation as before without any impact to the plant. However, with the cogeneration units operating at reduced load, this resulted in a reduction in steam generation and steam became a critical issue for the site. In order to generate more steam to meet the sites demand, one of the two cogeneration units was ramped up in load in order to achieve the minimum megawatt permissive for auxiliary firing of a Heat Recovery Steam Generator (HRSG). However, before this could be accomplished, the HRSG tripped on excessive superheated steam outlet temperature, which cross-tripped the gas turbine resulting in loss of the generator. With only one cogeneration system remaining on line and with no way to balance the site demand to internal generation, the demand exceeded the generation capacity resulting in a trip of the second unit. This dynamic event lead to a Site wide power outage lasting approximately 30 minutes, resulting in a loss of production, equipment damage, and an environmental impact.

It was recognized, after the second event, that if the power system could be dynamically monitored such that the internal generation and load could be balanced and, if required, pre-determined loads shed, then a stabilized island could be created and the plant electrical system could remain intact. Though the cost of installing an islanding and load shed system can be considerable, it was determined that the initial monetary capital cost of such a system would be eclipsed if a total loss of site power could be avoided.

Design Criteria

As mentioned earlier, the primary goal of the scheme is to detect an island condition and, if necessary, shed the amount of load required to create a load-generation balanced island. Step one is the detection of an island condition. From the diagram in figure 1, it can be seen that the "primary" island detection is accomplished by determining the Open/Close status of breakers MA and MB – the 34.5kV feeds into the plant. If both MA and MB breakers are open, then the plant is islanded from the main grid. A "secondary" island condition can be created if all four breakers in the plants 230kV ring bus are opened.

In addition to a "detected" island, the scheme was designed to "force" an island. Specifically, if an underfrequency or undervoltage condition is detected, the scheme trips both 34.5 Main breakers - MA and MB - to forcibly separate from the main grid due to an apparent unstable condition on the utility system.

Load Shed Decision Criteria

The decision as to whether to shed load and how much load to shed is based on the measurement of the dynamic load-generation balance. The internal plant load is calculated by summing the power flows on the 4 primary feeds into the plant, specifically:

$$Total\_Internal\_Load = P_{GA} + P_{MA} + P_{MB} + P_{GB}$$

Note that if there is excess generation from the internal generators, the power flow through breakers MA and MB becomes negative and the Total Internal Load is still calculated correctly.

The Total Internal Generation is then calculated as:

$$Total\_Internal\_Generation = P_{GA} + P_{GB}$$

The final calculation is the Load-Generation balance, which is:

$$Load\_Generation\_Balance = Total\_Internal\_Load - Total\_Internal\_Generation$$

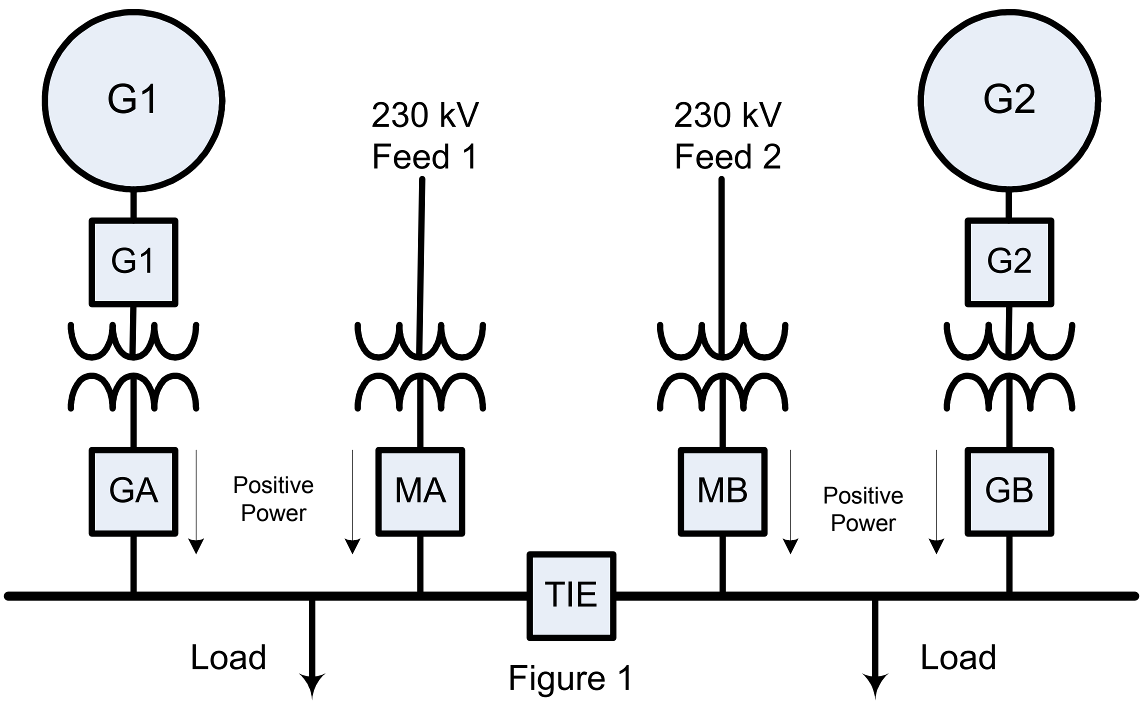

Figure 1
34.5kV Bus Sources

A positive value for the Load-Generation balance indicates that the load is greater than the available generation and that, upon detection of an island condition, a load-shed may be required in order to maintain plants electrical system stability.

The loads to be shed were identified by plant personnel and broken down into 3 load groups or “tiers”. For the specific design, the Load-Generation difference levels at which each tier was invoked were set at:

1.5 MW < (Load – Generation) ≤ 19.8 MW; Shed TIER 1 Load
19.8 MW < (Load – Generation) ≤ 27.3 MW; Shed TIER 1 and TIER 2 Load
27.3 MW < (Load – Generation); Shed TIER 1 and TIER 2 and TIER 3 Load

Note that the gas turbines were capable of picking up at least a 1.5MW Load-Generation difference and were quite fast at slowing down under the over-generation scenario.

Once islanded, there was still a chance that events in the plant could start to take down the local island. To address the "sinking island" scenario, and given that there was additional load to shed, two of the relays were programmed to address additional stability criteria, specifically:

If: Frequency < 58.8 Hz for 1.0 Second or

Voltage < .85 pu for 1.5 seconds

Note that is was necessary to coordinate the Underfrequency element on the generator protection relays with the islanded underfrequency load-shed values.

Depending on how many tiers of load were available to shed, different shed criteria were defined. To visualize the different states of the load shed system, a State Transition diagram was developed (figure 2) and used in the system design and as a system operating map and training tool.

As the distances between the controller and the sheddable loads were substantial (1000m), it was not practical to have direct copper paths to perform the load-shed trips. The decision was made to perform the trips via remotely located controllers. The signals to trip would be sent via a communication channel that would connect the various venues. As communications was crucial to the proper execution of the shed command, the decision was made to provide redundant communication channels. In addition, all communications were to be carried via fiber optic cable.

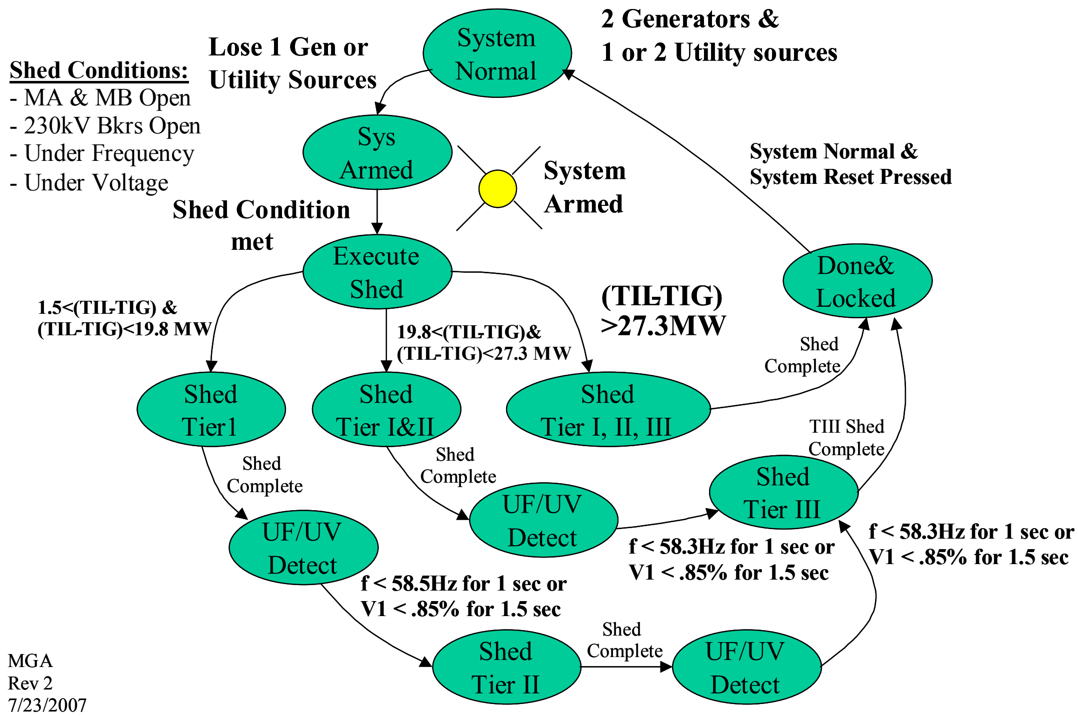

Figure 2
Load Shed Scheme State Transition Diagram

Implementation

The first part of the implementation was the calculation of the Load-Generation balance. To implement this calculation, programmable relays were located in the breaker cubicles of MA, MB, GA, and GB. These relays were connected to measure "positive" power flow as being into the plant. In addition, breaker status was monitored via digital inputs into these relays. As both "a" and "b" contacts were available from the breakers, double-point breaker status was implemented. In order for a breaker to be declared "open", the "a" contact had to be opened and the "b" contact had to be closed. If both the "a" and "b" contacts reported the same value, a "Breaker a/b Mismatch" alarm was issued. For reliability, a given breaker status was measured by two different relays. For example, the status of breaker MA was sensed by both the MA relay and also by the MB relay. The logic was designed such that a breaker was determined as "open" only if a ***valid*** open state was declared by either relay. Sharing of this information was accomplished through the use of both GOOSE (defined below) and a vendor specific Direct I/O communication scheme.

Load-Generation Balance

In order to perform the Load-Generation balance calculations noted above, all the analog data had to be communicated to a central location in the plant and then summed or differenced appropriately. The relays chosen for this application had a summator function available in the device. In this scheme, the MA relay was chosen as the calculation engine so all the analog values had to be communicated to the MA relay. The mechanism that was used to transport the analog values was the IEC 61850 Generic Object Oriented Substation Event or GOOSE. The GOOSE is a "user-defined" dataset that can be a combination of analog, digital, and quality data values. The GOOSE is launched on change of state of a status value, on a percent change in an analog value (user settable), or periodically as an integrity test. In this particular application, the analog value changes were tested once a second and re-transmitted as appropriate.

Once received and summed into the Load – Generation quantity, the resultant value was evaluated in comparators per the ranges noted in the Design section above. The output of the comparators was then mapped into "Tier" data values. Specifically, the Tier 1 Armed flag was set first, followed by the Tier 2 Armed flag, and finally followed by the Tier 3 Armed flag.

Load Shed Activation

By having the Tiers pre-armed, when an island condition is detected, the appropriate Tiers can be shed with no appreciable time delay. In the implemented scheme, the time delays come from 4 primary sources: Breaker Status indication, Logic Processing, Output Contact operation, and breaker operation. On the breaker status indication, all breaker status inputs are denounced for 8ms before becoming "valid" to use in the logic. Once validated, the breaker status values are fed into a logic engine that operated every 2ms. Any decisions to shed load were then communicated via GOOSE to the appropriate remotely located controllers. Note that the communication time is not noted as a primary source of delay as the time on the wire of the communication message is less than 300µsec. In the receiving relay, there is a 0 to 2ms delay for logic processing, there is a 3ms time delay in the operation of the output contact and finally, there is a 37ms breaker operate time. In as much as these times are "worst case" timings, an average time of 13ms was measured from initiation of the island condition to energization of the trip coils on the load-shed breakers or lockouts. Note that lockouts were needed on some loads to prevent automatic or operator re-starts.

Time Synchronization

The plant load-shedding scheme also included a GPS Satellite clock, used to time-synchronize each of the installed relays and controllers via different mechanisms and to varying degrees of accuracy. The relays used have synchrophasor recording capability, which is used as a long-time trend recorder in the operation of the plant electrical system. Recording is initiated via a number of different triggers. Time-synchronization of each of these synchrophasor relays required 1µsec absolute time accuracy. This level of accuracy was achieved through the use of the IRIG-B time synchronization protocol via direct wiring from the GPS clock to each relay. As the load controllers were located a significant distance from the GPS clock and as they did not need the same level of accuracy, time synchronization was achieved over the communication channel using the Simple Network Time Protocol – SNTP. Time synchronization via SNTP is typically able to achieve 1ms time accuracy. Given that all the devices are time synchronized, an integrated Sequence of Events report can be created that interleaves all the events from all the devices into one common report.

Device-to-Device Communications

Primary communications to all devices was achieved via Ethernet. As the operation of the system is based on reliable/available communications, a redundant Ethernet network was implemented. All the relays had redundant Ethernet ports on them and each port was connected through redundant Ethernet switches (see figure 3). The switches themselves were connected in a Ring configuration. In as much as Ethernet abhors a ring, an Ethernet protocol, known as Rapid Spanning Tree, was configured which automatically detects any ring conditions, dynamically breaks the ring, and automatically heals the ring upon detection of failure of another part of the ring.

In addition to the redundant Ethernet communications, a secondary communication path was established using a 64kbps synchronous communication channel connected in a ring amongst all the relays in the system. The secondary communication channel provides additional back-up communication of status and control information. As it operates in a ring configuration, there is a 3ms delay at each node as a received packet is received and the forwarded. This secondary channel was also used to communicate additional information that wouldn't fit in the available GOOSE packets.

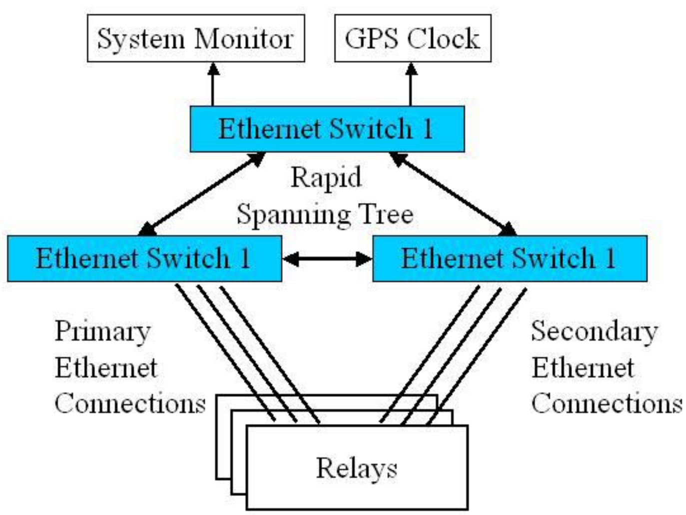

Figure 3
Redundant Ethernet Communications

User Interface

Monitoring and control of the system was achieved through LEDs and Pushbuttons located on the front of the relays. Specifically, one relay was chosen as a central controller that could arm/disarm all the devices in the system, RESET all targets and alarms, and control what breakers tripped in a given load shed scenario.

Status, operational information, and alarms of the system were visually reported through an assignment of the various pieces of information to user-programmable LEDs. Specifically, breaker status, arming levels, relay targets, device alarms, and other equipment alarms (e.g. – Ethernet switch status) were all mapped to individual LEDs on the front of the various relays.

Operational Information

In addition to the User Interface information, each of the relays provides additional operational information in the form of Sequence of Events (SOE), High-speed Oscillography, Synchrophasors, and Demand data. All digital inputs, internal logic status, and internal alarm information is captured in the SOE log. The contact inputs are sampled with a Δt of 500μsec and the internal logic and alarms are sampled with a Δt of 2ms. Each relay can hold up to 1024 events, however, an external system automatically retrieves the SOE logs from each device, integrates them into a common database, and provides standard database queries to sort and process the SOE information.

For events such as faults, high-speed oscillography is provided with a sample rate of 3840 samples per second. Long-term oscillography is provided by the synchrophasor function. Triggered by functions such as over or under frequency, over or under voltage, and rate of change of frequency, a single synchrophasor record can span over 20 minutes in length. Finally, 15 minute phase current, MW and MVar demand information is logged and can be retrieved for load analysis.

Testing

Testing of the system took place on two fronts, namely: a Factory Acceptance Test (FAT) and a Site Acceptance Test (SAT). For the FAT, the system was racked (all 8 relays), wired, and connected for communication (fiber cables to redundant Ethernet switches and Direct I/O) in the factory almost exactly as it would be installed in the field. The test included injection of voltages and currents into the measuring relays, simulation of all the contact inputs (representing the various breaker and lock-out states), and monitoring of the output contact performance. The GPS clock was connected and time synchronization was verified in all devices.

The voltage/current injection was wired such that the different tier arming scenarios could be simulated by opening, closing, and reversing the injected currents. As the design of the system was based on a state diagram, there was a clear test matrix that was derived based on tracing through all possible state transitions. Timing of the scheme was accomplished by examining the sequence of events logs from the relays involved in a particular Tier operation.

Site Acceptance Test

The SAT for the project consisted of two pieces: a commissioning test and a live islanding test.

The commissioning tests included:

- Verification of all contact inputs – as driven by the respective breaker or lockout relay. Note that this was greatly facilitated through the use of the relay front panel LEDs to which all breaker and lock-out status indications were mapped
- Verification of operation of each connected breaker or lock-out relay
- Verification of the various communication networks – including redundancy testing
- Verification of the time synchronization functionality
- Verification of power flow measurements and directions
- Verification of the various system alarms

Several minor issues were found during commissioning, however, the diagnostics designed into the system quickly identified the source of the issues and allowed for rapid remediation.

Of note in commissioning was the desire by the customer to change two of the breakers that were to be tripped during a load-shed operation. Due to the system design of tripping via GOOSE messaging, the changes involved writing some new logic and wiring the new breakers into one of the existing controllers – a relatively minor implementation task.

Live Islanding Tests

Once the system had been commissioned, the load-shed system was tested in two actual island situations – scenario 1 where a load shed was not required and scenario 2 where a load shed was required. As the plant was not operating at full load and as the available shedable load was only 3.5MW, the load-shed levels in the controller were temporality modified for a smaller shed range. The island was created by manually tripping the primary and secondary external feeds into the plant. As a precaution in both scenarios, the trip signals were removed from all but the available shedable load (a backup 3.5MW pump motor).

In scenario 1, the gas turbines were set to deliver the entire internal plant load plus an additional 9.4 MW, which was effectively exported to the local utility. Prior to the island, it was noted (as expected) that none of the load shed Tiers were "armed". Upon creation of the island, as expected, there was an immediate but small voltage increase (0.27%), however, within 16ms, the voltage regulator initiated a 0.87% drop in voltage (see figure 4). The turbines did start to accelerate but 1.78 seconds after islanding, the turbine controls had started the slow-down process of the turbine. The resulting positive sequence Synchrophasor voltage angle for the island is shown in figure 5.

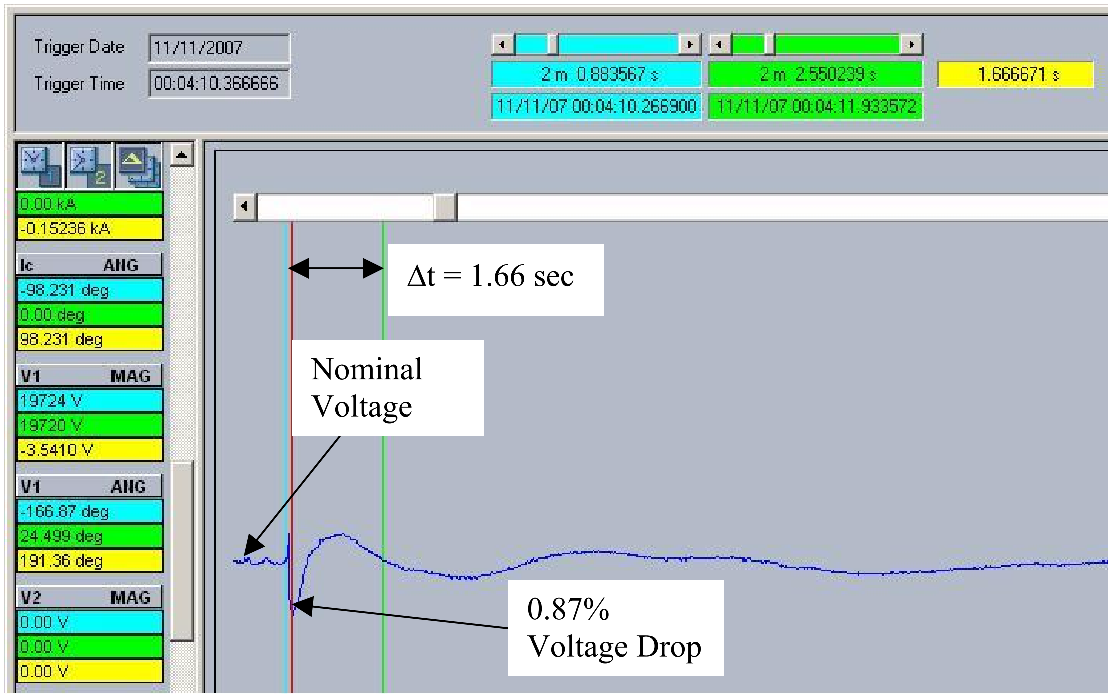

Figure 4

V1 Magnitude Synchrophasor Response to Over-powered Island Creation

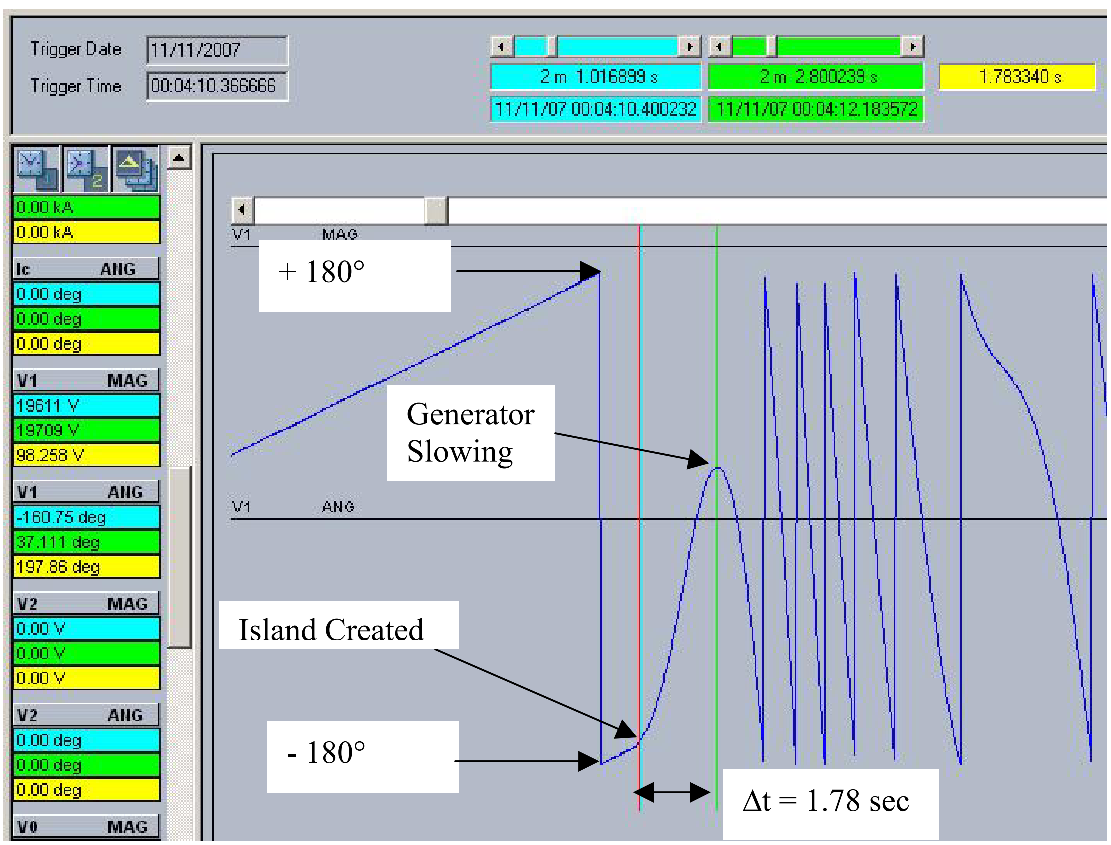

Figure 5

V1 Synchrophasor Angle Response to Over-powered Island Creation

In scenario 2, the turbines were programmed to output 4.55 MW ***less than*** the internal plant load. This scenario was designed to execute the "shed" commands. As mentioned earlier, the arming levels were temporarily lowered to force the arming of the respective load shed tiers. Prior to islanding, it was noted that all 3 Tiers were "armed" and ready to operate upon detection of the island condition. The recovery from the 4.55 MW deficit was to be made up of two sources, namely, a 3.5 MW motor shed and dynamic power up-take by the gas turbines.

Again, opening the main breakers created the island. The island was detected in 8ms (the debounce time on the breaker contacts) and a 3-Tier Load Shed command was immediately issued. The 3.5 MW motor was off-line (including breaker operation time) in 50ms from the detection of the island. Figure 6 shows the response of the positive sequence voltage – V1. Upon creation of the island, the voltage immediately drops only 0.21% and then starts to recover. When the motor load is shed, the voltage overshoots 0.66% - a very nominal amount.

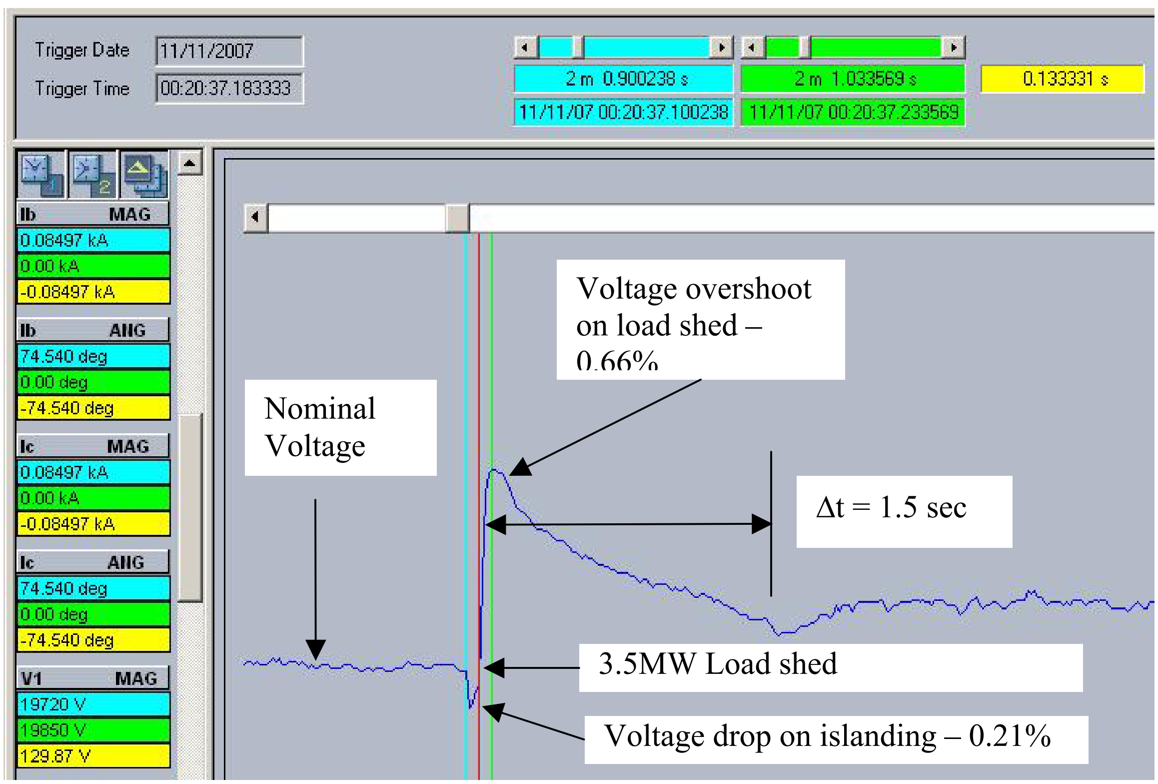

Figure 6
V1 Synchrophasor Magnitude Response to Underpower Island Creation

Conclusions

Today's manufacturing facilities require a higher degree of availability of electrical energy than in the past. Although load shed as a reliability mechanism is not a new concept, the design of a distributed system based on IEC 61850 GOOSE and an Ethernet communication network provides many advantages in terms of performance and flexibility as demonstrated through the inclusion of actual test results.