Considerations for Generation in an Islanded Operation

E. Roy Hamilton
Chevron ETC

John Undrill

Paul S. Hamer
Chevron Energy Technology Company

Scott Manson
Schweitzer Engineering Laboratories, Inc.

CONSIDERATIONS FOR GENERATION IN AN ISLANDED OPERATION

E. Roy Hamilton
Member, IEEE
Chevron ETC
1600 Smith Street
Houston, TX 77002, USA
erha@chevron.com

John Undrill
Fellow, IEEE
31 Via Maria
Scotia, NY 12302, USA
jundrill@ix.netcom.com

Paul S. Hamer
Fellow, IEEE
Chevron Energy Technology
Company
100 Chevron Way
Richmond, CA 94802, USA
paulhamer@chevron.com

Scott Manson
Member, IEEE
Schweitzer Engineering
Laboratories, Inc.
2350 NE Hopkins Court
Pullman, WA 99163, USA
scott_manson@selinc.com

Abstract—This paper discusses the conceptual design and operation of an isolated power system, recognizing the reality that generator or turbine trips will occur. The level of reserve generating capacity must be set with proper balancing of capital expenditures and operating costs against revenue lost in a production shutdown. The way that reserve capacity is provided is as important as the amount of reserve; seemingly adequate reserve can turn out to be badly insufficient if it is not well distributed across the available reserve sources.

The dynamic behavior of reserve capacity, as much as the amount of capacity that is ultimately available, is critical in determining how an isolated facility will behave in the wake of a unit trip or the loss of a grid connection.

In this paper, experiences with detailed dynamic simulations of a range of isolated systems are described. These are related to test work and operational incidents that have provided practical calibrations. Based on simulation and experience, some guidelines are offered for configuring generation and selecting strategies for maintaining stability in large, isolated continuous-process facilities.

Index Terms—dynamic stability, islanded power generation, incremental reserve margin, model validation, single-shaft gas turbine, spinning reserve, transient stability, load shedding.

I. INTRODUCTION

A common misconception is that the so-called "spinning reserve" in a power system can be evaluated satisfactorily by simply summing up the amount of connected generating capacity and subtracting the amount of connected load. An isolated (or islanded) system designed simply to have the difference between these totals greater than the largest potential loss of generation, or increase in load, is not very secure. The power system will experience difficulties in disturbances that seemingly should not affect it. This becomes extremely important when the facility's power system operates in isolation from a utility grid, either as a normal condition or in the wake of an event that interrupts a connection to a strong grid.

An islanded power system poses different operational "dynamics" on power generation units than those found on a strong utility grid. This paper focuses on the operation of large industrial-frame turbine-generation units in an islanded power system. It examines the limitations of turbine and governor response, the importance of accurately modeling the dynamic response of the turbine, validation of a turbine model, system design and operational considerations of multiple units in the islanded system, and the importance of a proper load-shedding system to ultimately maintain power system stability.

II. CASE STUDIES

This section relates experiences with two large industrial oil and gas production complexes for which the authors did extensive analytical studies.

A. Asia

The first complex is a large oil and gas production system located in Asia. This complex is capable of producing approximately 600,000 barrels of oil per day. It has a distributed power generation complex as shown in simplified form as Fig. 1.
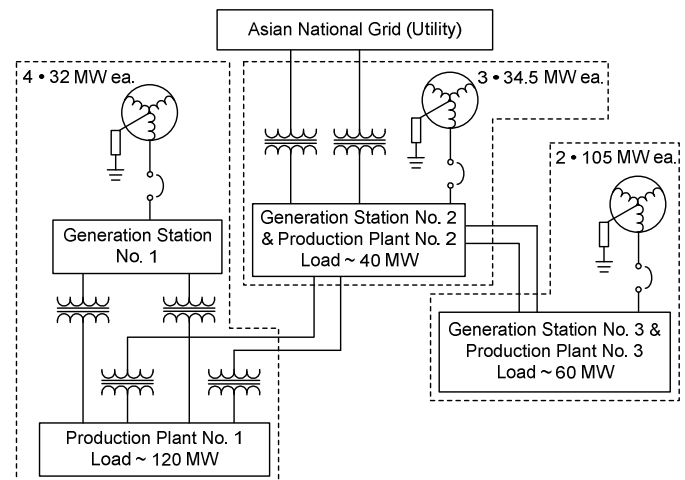

Fig. 1   Simplified One-Line Diagram of the Asian Complex

The system has three distinct production areas, each having a power generation station. Generation Station No. 1 has four early-generation, 32 MW, single-shaft, industrial-frame gas turbine units. Generation Station No. 2 has three relatively new, 34.5 MW, single-shaft, industrial-frame units. Generation Station No. 3 has two 105 MW units. All generating units are industrial-frame, single-shaft, gas turbine-driven, air-cooled generators. The three facilities are connected together by 110 kV redundant tie lines. Generation Stations No. 2 and 3 are outdoor, air-insulated substations arranged in a double-bus, single-breaker arrangement. A redundant, limited-capacity utility tie with the national grid is maintained for standby power import. It should be noted that with the power system arranged

in this manner, the 110 kV substation for Generation Station No. 2 becomes a power wheeling substation between Generation Station No. 3 and Generation Station No. 1.

Prior to the recent addition of Generation Station No. 3 and its associated production load, six of the seven units of Generation Stations No. 1 and 2 were used to handle the entire system power load. A spinning reserve margin of approximately 20 MW was left between the total generation and total load. Historically, if a single generation unit tripped, the power system was minimally impacted. The limited capacity tie to the utility and power that the remaining generation units could quickly assume, called the incremental reserve margin (IRM), easily picked up the load displaced by the tripped unit.

The two large units at Generation Station No. 3 were added with the expectation that these machines could be operated fully loaded and supply 75 percent of the total complex load. The remaining 25 percent of the complex load would be supplied by the three newer generation units at Generation Station No. 2. The Generation Station No. 1 units could be shut down and either dismantled or maintained as standby units held in ready reserve.

Neither the project personnel nor the power generation OEM (original equipment manufacturer) representatives understood the implications of having two generation units carry the majority of the power system load and the dynamics imparted by tripping one of these large units at full load. In the event of the loss of one large generation unit, the power system would now be required to pick up 37 percent of the load on the remaining online machines; whereas in the past, the loss of a single, smaller unit represented a pickup of approximately 19 percent of the load. The sudden step of 37 percent of total system load on the online machines represents a significant event on this power system, even with the assistance of the limited utility tie. Study work showed that the response capability of the remaining online generation units was not sufficient to handle such an event. The utility tie helped to provide immediate incremental reserve, but study results indicated varying amounts of load shedding might still be required to maintain the system.

B. Indonesia

The second case study involved a large oil production complex located in Indonesia. This complex produces approximately 300,000 barrels of oil per day and has a distributed power generation complex as shown in simplified form as Fig. 2.
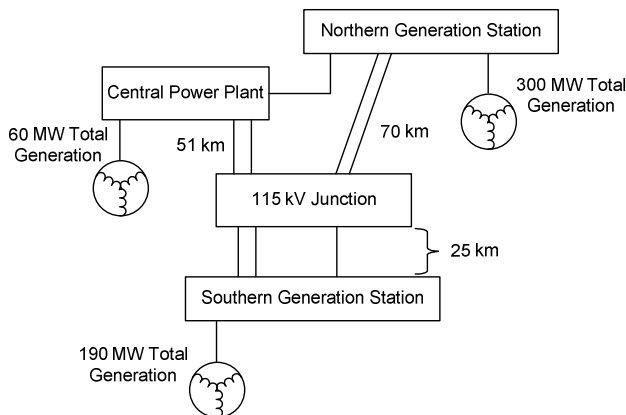

Fig. 2   Simplified One-Line Diagram for Indonesian Islanded Power System

Power generation for this complex is essentially located in two areas. The older portion of the facility has a southern power station with eight 20 MW units and three 35 MW units. A newer northern generation station has three 105 MW units. All generating units are single-shaft, industrial-frame machines. The entire power system is self-contained (islanded) with no connection to the external national grid.

The northern generation station produces both power and process steam and is located over 70 kilometers away from the southern generation station. The two generation stations are interconnected by a 230 kV transmission line. Power is distributed to the entire production field by 115 kV transmission lines. The loads on this system are predominantly induction motors. The total system load averages 430 MW. To support this load, the three large generation units at the northern station are operated at nearly base-loaded condition, and generally, eight smaller units at the southern station are operated with a spinning reserve margin of approximately 40 MW.

This facility experienced an instrumentation failure on one of the large northern station turbines. The result was the tripping of the unit and, within minutes, the collapse of the entire power system. Multiple layers of underfrequency load shedding totaling more than the lost generation were triggered, but the system still collapsed.

In the case of the Asian system, a planning study anticipated operational problems. In the Indonesian case, studies were undertaken to explain the behavior of the system after the event and to plan measures to prevent recurrences. A review of the Indonesian event revealed that the magnitude of the load-shedding stages initiated by underfrequency relaying was too small to arrest the frequency decay. This is further explained in the next section.

III. LIMITATIONS OF TURBINE GOVERNOR RESPONSE

So why was an underfrequency load-shedding system with what was thought to be adequate spinning reserve unable to save the Indonesian system? Was it not possible for the turbine governors to simply push the turbines to utilize their spinning reserve?

The governor controls for large, industrial-frame turbine generators have traditionally been designed with the expectation that they will be connected to a strong utility grid and that speed (frequency) variations will be minimal. The oil and gas industries, however, are placing facilities in remote locations where there is no utility tie or the utility system is weak. In these environments, the turbines can experience large variations of speed, and their ability to respond to these changes is critical to the security of the system. The events may be as simple as a large motor start, loss of a generation unit, failure of a switchgear bus or transmission/distribution line, or loss of the fuel gas supply to the engines. Understanding the response limitations of the engine is very important.

It is also important to understand that the arithmetic difference between connected capacity and load is not a useful indication of the generating units' ability to pick up load on the time scale needed to arrest a rapid fall of frequency. Terminology is important. This paper uses "spinning reserve" to refer to the difference between present turbine-generator output and the maximum that can be achieved, up to thermal limits, given sufficient time, without starting another unit. Additionally, "incremental reserve margin" is used in this paper

to describe the increase in output that can be achieved in a short time interval, typically 5 to 10 seconds.

As a general guide, the gas turbine can be expected to quickly pick up 10 to 15 percent of its site-rated capability and then assume additional loading at a rate of 0.5 percent (of its site output rating) per second. This is very different from assuming that the full spinning reserve is instantaneously available.

Fig. 3 illustrates this delayed response. The Indonesian 100 MW machine was loaded to 92 MW, and a step load of approximately 5 MW was added at 5 seconds. The solid black line (Series 1) represents turbine speed (or system frequency). Notice that it dips as the step load is added. The thin black line (Series 3) represents the electrical power output of the generator. The change in electrical output is instantaneous. The gray line (Series 2) represents the mechanical power output of the turbine. It is not an instantaneous step. The initial instantaneous electrical change is determined by Kirchoff's laws and system impedance; it is not affected by turbine characteristics or even by turbine-generator inertia. The subsequent oscillatory component of electrical response is determined by impedances and mechanical inertias. The response of turbine power, on the other hand, is determined by the combined influences of the thermal characteristics, the turbine governor, and the supervising elements of the turbine controls that are intended to avoid damaging conditions in the turbine.

There is a clear lag between the generator output and the turbine power output. It took nearly 5 seconds for the turbine governor and the fuel control valve to add enough fuel to increase the turbine power output to match the power demand on the generator.
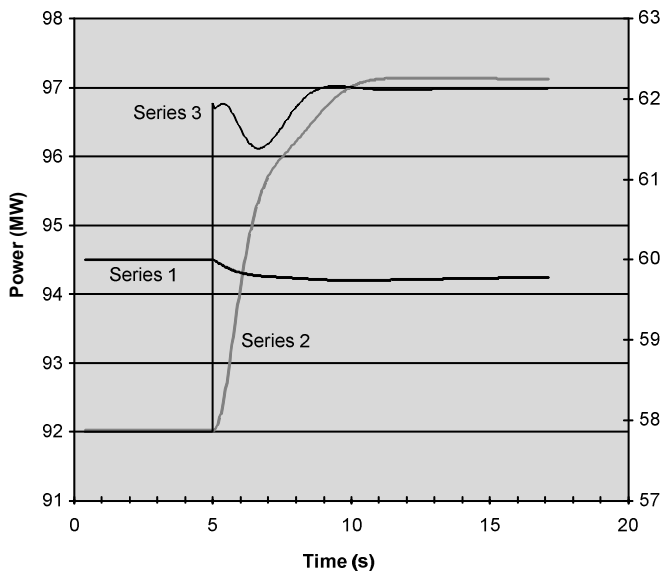
The decay of speed can only be arrested and corrected when, and if, the governors can bring the collective turbine power up to match and exceed the collective load. In this context, it is important to recognize that the maximum output allowed by the gas turbine controls is reduced roughly in proportion to the square of the speed as the turbine speed decreases. This limitation on output does not appear instantaneously; it is imposed as the measured exhaust temperature rises and the temperature limiting controller takes command of fuel flow. The limitation on turbine output may cause an event that seems initially to be survivable to evolve into a power system collapse.

In Fig. 4, three units of the isolated system were running near site-rated base load, supplying a total load of approximately 280 MW. A unit carrying 93 MW tripped, causing a very rapid decay of system frequency. Significant load shedding was necessary because of this sudden loss of 33 percent of online generation on essentially base-loaded machines. Underfrequency relaying was set to drop load in excess of equivalent generation in several stages of approximately 20 MW each. The dashed line (Series 1) represents turbine speed (or system frequency). The black solid line (Series 3) represents the electrical power output of the generator. The gray solid line (Series 2) represents the mechanical power output of the turbine. Again, the electrical load on the generator jumps instantly. The turbine governor starts to respond but is limited immediately by the temperature limit. Speed decreases rapidly under this large step load. Underfrequency relaying operates, and at about 12 seconds, speed appears to stabilize. At about 16 seconds, the system seemingly recovers. Fig. 5 shows this same exact plot, but with time extended to 24 seconds. At 18 seconds, it becomes obvious that the system is crashing.

Fig. 3   Power Output and Speed of an Industrial-Frame Turbine and Generator, Small Load Step Response
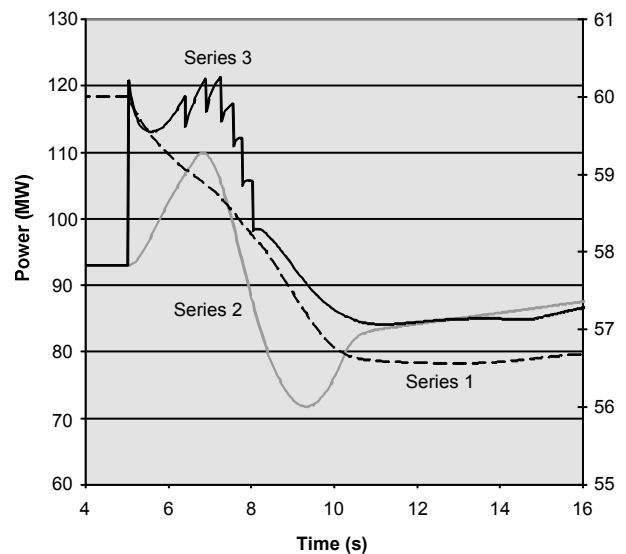
Fig. 4   Power Output and Speed of an Industrial-Frame Turbine and Generator, Apparent Successful Recovery

The difference between electrical power and turbine power creates an accelerating or decelerating torque on the turbine-generator shafts. The system speed (frequency) continues to decay until the collective turbine power output matches or exceeds the collective electrical demand. In this particular instance, speed dipped by only 16 rpm before recovering.
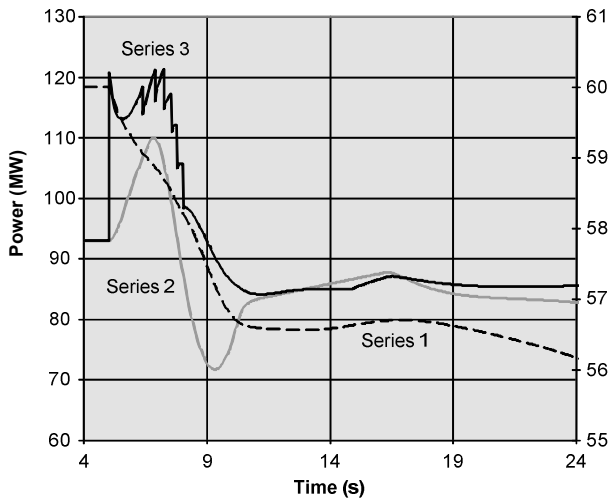
Fig. 5   Power Output and Speed of an Industrial-Frame Turbine and Generator, Failed Recovery

Studies of system behavior must accurately take into account turbine limitation such as is shown in Fig. 4 and Fig. 5. This requires accurate models of governors and turbine controls for valid simulations.

When a given facility load is covered by a large number of generating units (e.g., five or more), the loss of one unit might be covered adequately by the dynamic response of those remaining online. If a smaller number of larger units are used

to cover the load, the dynamic response of those remaining will not be able to cover the loss of one unit, even if the arithmetic summations suggest that there is adequate spinning reserve. Under these circumstances, load shedding will be essential.

Capital planning of unit sizes, the design of load-shedding systems, and the broad range of related operational issues all require dynamic simulation studies in which the response capabilities of turbines are modeled comprehensively.

IV. VALIDATING THE TURBINE AND TURBINE GOVERNOR MODEL

Effective studies of reserve response require both computer programs that provide the appropriate dynamic models and assurance that these mathematical models are properly calibrated. The presence of a model in the library of a computer program does not give the required assurance of proper representation. Models must be calibrated against the measured behavior of in-service units.

A typical model is shown in Fig. 6; this example represents a 120 MW industrial-frame turbine and its principal controls. This model's form was reviewed by comparison with site-specific, as-built control schematics. As-built control information provides reliable values for some parameters, such as droop setting, but cannot provide calibration with regard to the behavior of the turbine itself. Overall calibration is best achieved by comparing simulations with the recorded results of response tests.
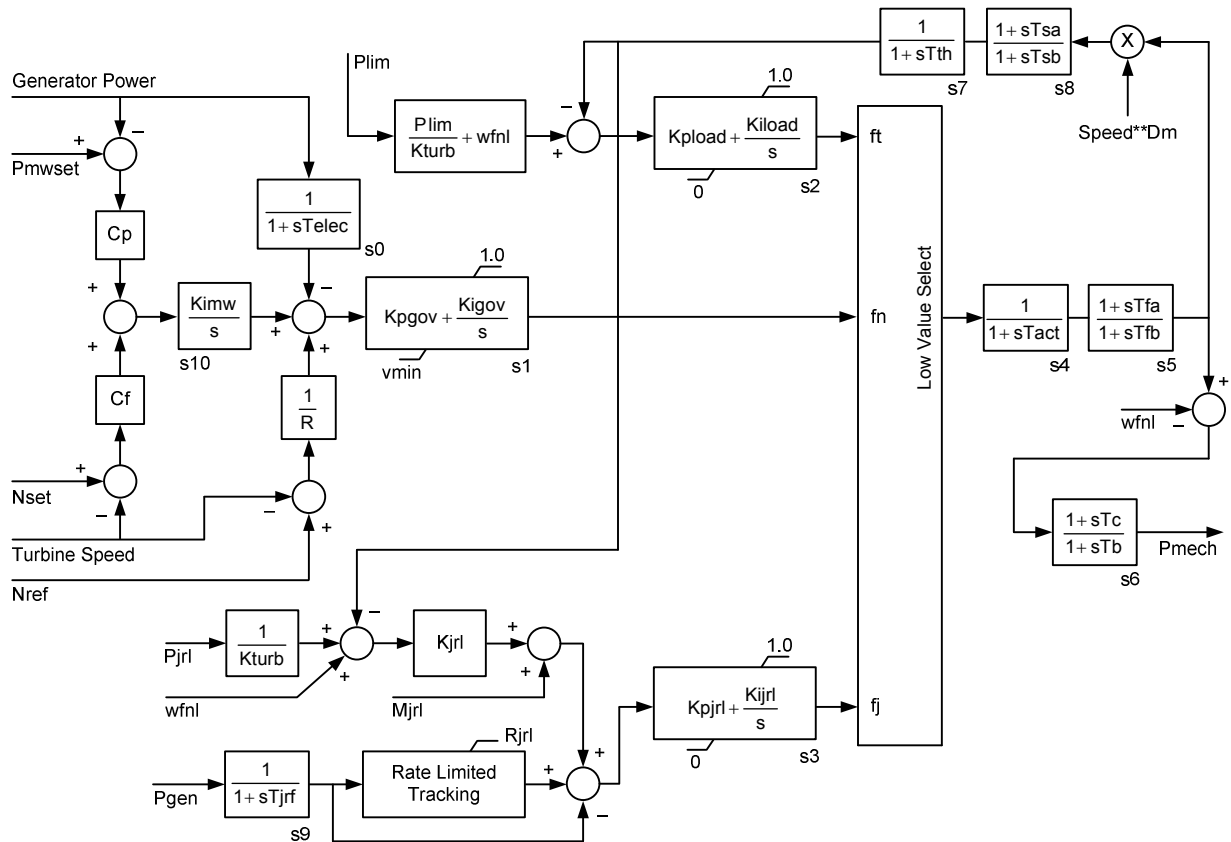

Fig. 6   Generic Governor and Turbine Model (WGOV1)

4

A. Tests for Model Validation

The response tests needed for validation of dynamic performance models can be undertaken generally as follows:

1. Test a unit that will fairly represent the type under consideration.

2. Record the following signals at a rate of at least 10 samples per second:
   a. Generator MW
   b. Turbine fuel command
   c. Turbine compressor speed (if a multishaft engine and if available)
   d. Turbine power shaft speed (generator speed or frequency)
   e. Turbine exhaust temperature
   f. Turbine pressure ratio (if available)

3. If circumstances allow operation into a resistive load bank, switch the load on and off to produce steps ranging from 5 to 15 percent of the rated output.

4. Where the generator must be tested in its service connection, initiate changes of output by making step changes to the governor speed-load reference so as to change the output by amounts between 5 and 15 percent. Note that the output changes of the test may change the frequency of the power system if the test generator is large in relation to the system of which it is a part.

5. In both test situations, apply output increase steps that will take the turbine decisively up to its exhaust temperature limit; for example, apply a step that would take the turbine from 95 to 102 percent output in the absence of the limit.

It can be anticipated that test personnel would need to be on site at least a day prior to the test to discuss the test procedure, prepare operational loading plans, and set up recording systems. It is imperative that the recording system be tested and proven in normal operation prior to the start of testing.

B. Model Validation

Testing does not validate a dynamic model. Model validation is an analytical process based on test results. Each of the tests made as described above should be simulated with the dynamic model proposed to represent the machine. The parameters of the model are adjusted until the behavior shown by the model in a simulation of each test is a fair match to the observed test behavior. A perfect fit between a single test recording and simulation result is rarely achieved and is less important than achieving a fair fit over a range of magnitudes of test disturbance and a range of initial loading levels. This process requires exact knowledge of the way the tests were conducted, expertise in modeling, and sound knowledge of the turbine and controls under consideration. This analytical part of the exercise is more demanding than the test phase in terms of the availability of expertise, elapsed time, and cost. Attempts to minimize cost by having test work done by field technical staff and analytical work by other specialists have been notably unsuccessful.

It is fair to regard dynamic model validation as a significant expense when taken on its own. However, the costs of failed system performance because of inaccurate analysis in the planning stage are significantly greater.

Fig. 7 shows the measured data from a step response test of one of the engines located at the Indonesian facility. The three 100 MW units were connected to an islanded utility grid and loaded to approximately 85 percent. The governor speed-load reference for northern generation station Unit No. 2 was stepped upward to make the engine increase its output by 8 MW. After 50 seconds, the speed-load reference was stepped back to its original setting. The plot shows the fuel command signal, rotor speed, (electrical) power output, and power reference signal. Because of the large size of the turbine in relation to the size of the isolated power system, it was necessary to limit the test steps to 8 MW; the turbine could have made larger steps, but these would have caused unacceptable changes in system frequency.
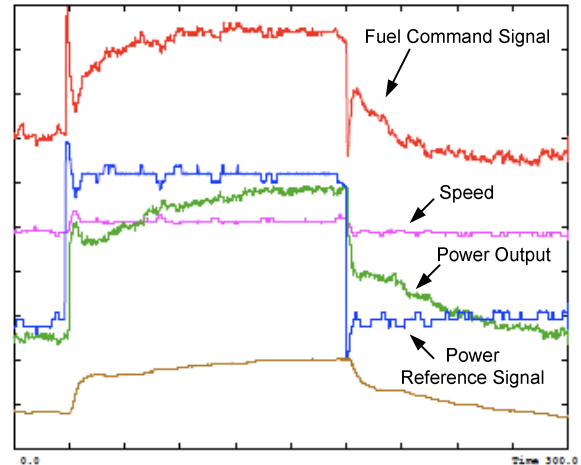
Fig. 7   Actual Plot of Test Data for a 110 MW Industrial-Frame Turbine Engine

Fig. 8 shows the model simulation of the test shown in Fig. 7 after all model parameters had been adjusted to correspond to known as-built values where applicable and to give the required fair match of simulation to test. The good correspondence between simulation and test validates the model for use in system studies.

Fig. 8   Plot of Results for a Model of the Engine Plotted in Fig. 7

V. SIMULATION STUDIES AND GENERATOR CONFIGURATION

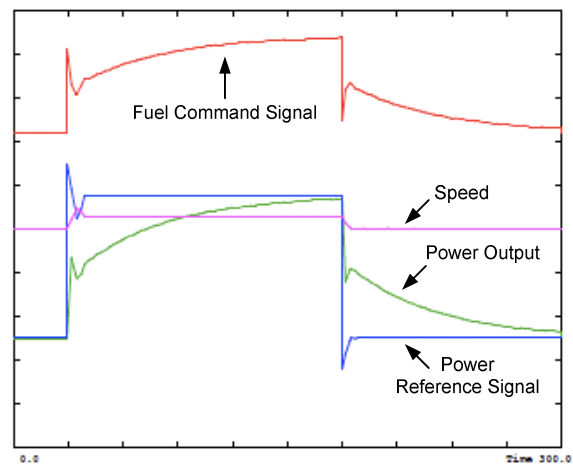What should be the approach to validating the configuration for an islanded power generation system? The first step is to realize that there is a difference between an electrical transient stability study and a power system dynamic study. The transient stability study usually concentrates on events that are of very short duration, such as power system short circuits and clearance by circuit breakers, and whose time window of interest is up to approximately 1 second from the event's appearance. This time duration is too short for the turbine governors to respond. A dynamic study looks at the effects to the power system for approximately 1 to 20+ seconds following the event.

The second step is to develop an accurate system model for the facility's power generation and power distribution (including the facility loads). This model can then be employed for transient and dynamic system response studies with the user looking very carefully at the behavior of the power generation elements. For the transient stability study, a grossly simplified model of a governor is acceptable or may not even be required. However, a system dynamic study requires a governor model that contains accurate depictions of the thermal, mechanical, and governor limits of the turbine. A system can seemingly survive a transient disturbance only to collapse some time later because of the protective controller actions from the governor.

The third step is to subject the system to all possible contingencies it may see. This includes the loss of one generation unit, a fault on the generator bus, a fault on a distribution bus, a fault on a cable feed, the start of a large motor, the loss of a bus coupler or tie line, and the loss of significant load.

The fourth step is to analyze the simulation results and validate them as credible. This means that the user has to understand the capabilities of the equipment being modeled, which in turn implies the user must have a certain level of experience with this equipment.

The fifth step, once the system dynamics have been properly modeled and understood, is to properly implement protection systems that will take action and ensure that power system stability is maintained.

VI. LOAD-SHEDDING PROTECTION SYSTEMS

Power system collapses quite commonly occur because the power system frequency decays at an extreme rate and protection systems trip off motors and generators, causing further generation to trip, eventually cascading into a full system outage. However common the outcome, the origination of a system outage can have many different initiating factors. The initiating events for a system outage may have occurred seconds, minutes, or hours prior to the collapse. For small industrial and islanded power systems, the most common form of initiating event is the sudden loss (circuit breaker trip) of a generator, bus coupler breaker, or tie breaker. If any of these breakers suddenly are opened (under load), a power imbalance will occur between the mechanical power created by the turbines and the net sum of the electrical load on the power system. This section deals specifically with proven remediation methods used to rebalance the remaining turbines and loads, thereby preventing system frequency decay.

The frequency decay rate of an electrical system under a power deficit is related to the magnitude of the power deficit, the load composition (induction motor, synchronous motor, resistive loads, electronic loads), and system inertia (H constant). For approximately the first second, this decay will occur regardless of the type or quality of the turbine governor. As an example of both extremes of governor control action, Fig. 9 shows the initial decay rate is identical for turbine governors running in base load (no speed control) and governors in droop mode (speed control with power bias factor). Therefore, for all power systems, an underfrequency load-shedding system will only detect a frequency decay after the initiating condition of a power deficit. As shown previously, this delayed response time can quite frequently result in a cascading blackout.
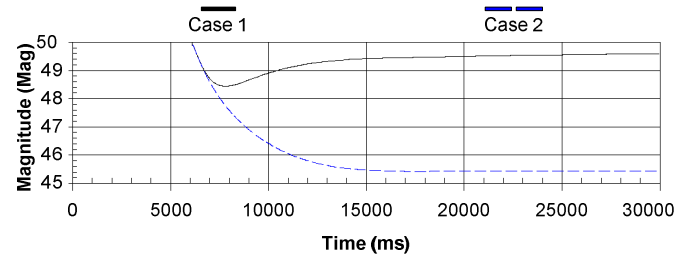

Fig. 9   Frequency Response Comparison of a Governor in Droop and Base Modes

A. Contingency-Based Load Shedding

A proven method for correction of power generation versus load unbalance is to shed (trip) loads immediately upon the opening of breakers through which active power is flowing. The opening of a generator, bus coupler, or tie breaker under load can create a power disparity. These are therefore classified as contingencies that can cause power imbalance. Therefore, this form of system is referred to as a contingency-based load-shedding protection system.

The contingency-based load-shedding protection system has many names and acronyms throughout the world. These systems are commonly called "special protection schemes" (SPS) or "remedial action schemes" (RAS) by electrical utilities. For industrial and commercial electric power systems, these protection schemes are most commonly integrated into an overall electrical power system protection package containing many hundreds of multifunction protective relays. These load- and generation-shedding protection schemes are commonly included in many industrial power demand management systems (PDMS).

Depending on the communications protocols and media, modern, contingency-based load- and generation-shedding protection systems can have closed-loop response times of faster than 12 milliseconds over hundreds of kilometers, thousands of contingencies, and tens of thousands of loads [1]. This time is the measured total time from an input voltage asserting to 90 percent of full voltage to an output contact fully conducting on a controller's I/O terminal blocks. This includes the full conduction of output contacts that are rated for tripping; therefore, interposing relays are no longer used in modern systems. Because of these speeds, contingency-based protection systems are now realistic for any size or type of power system.

Various signals have been used over the years to initiate a load-shedding contingency. These include breaker contacts (52 A and B contacts), 86 lockout contacts, current thresholds, out-of-step (OOS) conditions, protective relaying trip signals, synchrophasor phase angle deflection [2], thermal limits on

generators, transformer overloads, voltage depressions, and more. Each of the aforementioned contingency triggering conditions has an impact on the overall system-shedding time and the operational security of the overall scheme.

A single failure of a communications processing, logic processing, or I/O device can be catastrophic in a contingency-based protection scheme. It is for this reason that modern load- and generation-shedding schemes are built exclusively on protection-class equipment with substation environmental ratings. All protocols used are encrypted, protection-class signals communicated over devoted communications channels. For the most rigorous of applications, triple modular redundant (TMR) voting schemes are used [3].

Modern, contingency-based load-shedding protection systems perform all of their calculations and subsequent load selections continuously and prior to any contingency event [4]. In this way, the system is always armed with the appropriate load-shedding solution and continuously reports to the operators the outcome of every possible future contingency event.

The basic underlying equation used to select the amount of load to be shed is:

$$L_n = P_n - \sum_{g=1}^{m} IRM_{ng} \qquad (1)$$

where:

n = contingency (event) number

m = number of sources (generators) in the system

g = generator number, 1 through m

L_n = amount of load selected for "n" event (MW)

P_n = power disparity caused by "n" event (MW)

IRM_{ng} = incremental reserve margin of all generators (sources) remaining after "n" event (MW)

There are several key characteristics of modern load-shedding systems, including:

1. Pre-armed load-shed events, per (1). These arming signals are commonly loaded into a construct called a crosspoint switch matrix for ease of indexing and operator display.

2. Operator selection of sheddable load priorities.

3. Operator selection of IRM for each power source (generator).

4. Event logs (event reports) that capture detailed analog and digital information of each event that occurs, with up to 1-millisecond accuracy and time durations of up to 30 seconds.

5. Sequence of event (SOE) logs, which capture all changes of state of digital signals with 1-millisecond accuracy.

6. 1-millisecond or better accurate time synchronization of all electronics to coordinated universal time (UTC). This is most commonly accomplished by synchronization of all electronics with IRIG-B satellite time-synchronization signals.

7. System diagnostic logs to capture and time-stamp any equipment anomalies.

8. Real-time, temperature-compensated modeling of the long-term reserve margin capabilities of generators and turbines. This is used to provide realistic limits to any operator-entered IRM values.

9. System topology tracking. This includes complete tracking of all breakers and disconnect statuses carrying power between sources and loads. Load-shedding algorithms must know the routes in which power is flowing between sheddable loads and sources.

The inclusion of these basic concepts into a contingency-based load-shedding system is the reason many systems are described with such terms as "predictive," "flexible," "adaptive," or "intelligent." Reference [5] identifies a large number of other critical characteristics of these systems.

Modern, contingency-based load- and generation-shedding systems must handle multiple, closely timed events. Unfortunately, current and voltage values commonly oscillate following such a power system disturbance (contingency event). These transient oscillations are easily measured with modern electronics; however, without steady-state information, the evaluation of (1) becomes impractical. Allowing a contingency-based system to shed load based upon transient information will commonly undershed or overshed, possibly making a bad situation worse. Contingency-based systems are therefore commonly inhibited from tripping action for some time period following the first disturbance (contingency). Two methods are commonly employed to provide load-shedding protection for multiple, closely timed events:

- Queuing of contingency events and submillisecond power-flow recalculation. This power flow is used to determine the new steady-state conditions during times of transient oscillations. Such schemes are most commonly employed on mission-critical generation-shedding schemes for utility systems [3].
- Backup underfrequency load-shedding system. These systems provide protection for power disparities during the contingency-based system transient inhibit [4].

B. Underfrequency Load Shedding

Underfrequency load-shedding schemes are commonly employed in industrial power systems as a backup to a contingency-based load-shedding system. In addition to transient inhibit periods, maintenance issues such as equipment failures, broken wiring, shorted CT windings, and dc battery failures will cause a contingency-based load-shedding protection system to not operate when needed. Clamping and slew rate limiters in governors or fuel/air problems are other situations for which a contingency-based load-shedding protection system will not operate. Improper installation or commissioning of protection equipment can also be another reason that a contingency-based system will not react when needed. All of these reasons make it mandatory that a backup underfrequency-based load-shedding system be employed to supplement a contingency-based system.

Unfortunately, there are severe limitations in traditional underfrequency load-shedding protection systems, primarily because this type of system only reacts after the system is in a state of decay due to overload. These limitations have caused

load-shedding systems to gain a bad reputation as "untrustworthy." It is the authors' experience that underfrequency-based systems based on single-function underfrequency relays have an approximately 50 percent likelihood of rescuing a power system from decay.

Hybrid underfrequency load-shedding systems serve to correct all of the known weak points of traditional underfrequency systems. These hybrid schemes still shed load based on several underfrequency thresholds; however, the signals are sent from remote devices to a centralized processor. These protection schemes then dynamically select from a prioritized load list. It is this similarity to contingency-based systems that gives this category of underfrequency load-shedding systems the name "hybrid."

Table I summarizes the advantages of hybrid underfrequency schemes over traditional schemes that use underfrequency elements in remote, separate relays.

TABLE I
UNDERFREQUENCY SCHEME COMPARISON

Item	Hybrid Underfrequency Scheme	Relay Underfrequency Scheme
1	Selects correct amount of load to shed for every underfrequency threshold based on live power (MW) and knowledge of the power system R value.	Underfrequency elements operate with any amount of load through the shed breaker. May not shed load if load is off, or may shed too much if load is larger than anticipated.
2	Always sheds the optimal amount of load (MW).	Basically a fixed, nonadaptive system.
3	Sheds less load with better impact.	Commonly sheds too much load, sometimes resulting in power system instability or overfrequency.
4	Changing priority of sheddable load is very easy; just change the load priority from the user interface.	Changing priority requires changing underfrequency pickups and timers on discrete relays, very labor intensive.
5	No maintenance.	Regular maintenance and testing required on old single-function underfrequency relays.
6	Typically > 99.99999% of availability.	Typically < 99.99% of availability.

The advances over traditional underfrequency load shedding that these hybrid systems provide include the following:

- Loads are dynamically selected (only active loads are selected to be shed).
- Load consumption (MW) is incorporated into the selection of load to shed.
- Power system topology is tracked, guaranteeing that all loads that are shed are on the bus or island that required load shedding.
- The incremental change in frequency (F) versus power consumption (MW) is selected by the user ($\Delta F/\Delta MW$). This ratio is normally determined by a power system dynamic stability study.

C. System Modeling and Validation

It is imperative to characterize a power system before the configuration of any modern contingency- or underfrequency-based load-shedding scheme. This characterization is only possible with a hardware and/or software package capable of providing accurate "power system dynamic studies." This is not to be confused with "electrical transient stability studies," as previously discussed in this paper. The package used for load-shedding characterization must include accurate modeling of governors, turbines, exciters, rotating machinery inertia, load mechanical and electrical characteristics, electrical component impedances, and magnetic saturation of electrical components. Several of the parameters that come out of these dynamic stability studies include: the IRM of each generator and connected utility, ($\Delta F/\Delta MW$) of the system, coordination validation between underfrequency backup systems and contingency load-shedding systems, load makeup ratios, and total system inertia (H). These parameters are crucial to the proper operation and coordination of modern load-shedding systems.

Various levels of testing go into the validation of any complete load- or generation-based shedding. For the most mission-critical applications, a live real-time simulation environment is used to validate complete system performance. This is accomplished with a hardware package capable of providing both power system dynamic and electrical transient frequencies of responses (i.e., it must accurately model all mechanical and electrical systems). Such tasks are only accomplished with several hundred parallel processors running real time in a single-purpose simulation environment. To fully validate the load- or generation-shedding protection systems, the processors have direct, hard-wired connections to the protection equipment's CT, PT, and I/O connections. The authors' experience is that this form of real-time, closed-loop simulation is essential to fully validate any new generation of contingency-based load- and/or generation-shedding equipment.

VII. OPERATIONAL CONSIDERATIONS

The purpose of this section is to provide some guidance for applying the information provided in the previous sections of this paper.

The first consideration in selecting the proper generation configuration and load-shedding protection scheme is to identify the sensitivity of the process to power outage events. If the process can easily withstand significant loss of load without adverse safety implications to the plant and personnel and if lost profit opportunities are not a concern, the number of generation units can be closely matched to the operating load and loads shed as appropriate to minimize the outage. This type of situation may be found in an islanded power system, such as that found on floating production, storage, and offloading vessels called FPSOs. A significant amount of load that can be shed without adversely affecting the vessel's production may include water injection and sea water pumps. These may be of enough magnitude to compensate for the loss of one generation unit. It must be noted that even though the process can tolerate load shedding, the generation units must never be operated near the turbine's firing temperature limit. There has to be some margin left between the turbine's loading and the site-rated thermal maximum limit to allow for

small deviations in load. This can also be referred to as a margin for frequency maintenance. As small loads are switched on and off onto a small group (four or less) of base-loaded machines, the frequency of the system may start a slow decay. This becomes especially true if the small additional loads stay on continuously. Unless safeguards are in effect to automatically shed load or alert the operator to shed load, the frequency may decay over several hours to a point that one or more units reach a stall condition. The result will be a trip of the turbine to protect it from entering a surge condition. The other generation units are already at maximum loading when the displaced load of a tripped generation unit is imposed upon them. The whole system will collapse. Fast load shedding may not be enough to prevent the system collapse for a system operated in this manner.

If the process is deemed critical and load shedding should be kept to a minimum, another philosophy must be adopted. The authors recommend that for a critical process, the size of the largest generation unit should not exceed 20 percent of the total generation capacity. For a system of equally sized machines, this translates to a minimum of five equally sized units. In the authors' experience, for such a system with five equally sized machines operating with 20 percent thermal margin left on each machine, the loss of one unit is compensated for by the remaining four machines. The thermal margin will of course vary depending upon the ambient temperature, condition of the machines, and degradation of the air inlet filters. With less thermal margin, a minimal amount of load shedding may be needed to maintain the system's integrity. The optimal number and type of machines must be considered from an economic perspective as well. Factors such as capital equipment cost, maintenance intervals, and, in some instances, size and weight may influence the final selection of size, number, and type of turbine generation units.

Many islanded systems employ a large number of generation units connected to a generation bus or a main power distribution bus. In such situations, the user is encouraged to look carefully at the X/R ratio of the current imposed on the bus and circuit breakers during short-circuit conditions. It may be that the dc offset created by the higher X/R ratios of close-coupled generation may exceed the capabilities of the circuit breakers. Higher-rated equipment, delayed breaker operation, or series impedance may be needed to compensate for this condition and allow the equipment to operate safely.

Should the system architecture be designed such that the overall power system can be broken into separate islands, care must be taken not to separate load from generation by high impedance. This is especially true if an expansion to an existing facility may be adding additional load and generation that would be coupled to the existing system by tie lines or transformers. If too much impedance is required to connect the various system parts together, a bottleneck to reactive power flow will be created, and voltage support issues will arise.

A power management system with automatic generation control may be considered to operate a system with multiple generation units. This system can monitor the topology of the system and dispatch the generation as needed to load each generation unit in equal percentage. Reactive power flows can be adjusted by this type of controller through transformer on-load tap changers. Fast load shedding and fast generation shedding can also be implemented as a subset of this system but should be kept separate for system integrity reasons. Both fast load-shedding and fast generation-shedding systems need

to maintain optimal response times less than 200 milliseconds (including breaker opening times) to be effective. Times longer than 250 milliseconds enter the time-response capability of traditional underfrequency relaying.

VIII. CONCLUSION

The dynamic response of the turbine governor becomes particularly important when the engines are operated to their limits. This can be under normal operation of base-loading the engines; it can be through imposition of large step-load demands on the engines through loss of generation or through the addition of significant blocks of load.

Islanded power systems with fewer, larger generators must rely more heavily on strategies like contingency-based load shedding. Traditional, underfrequency-based load-shedding systems are not appropriate as primary blackout remediation techniques for islanded systems. Hybrid underfrequency schemes have proven to be an appropriate backup scheme if they are properly coordinated with a primary, high-speed, contingency-based system. Preferably, all load-shedding system coordination and controls must be validated with a full power system dynamic study with generator models that have been validated against real data.

IX. REFERENCES

[1] B. Cho, M. Almulla, H. Kim, and N. Seeley, "The Application of a Redundant Load-Shedding System for Islanded Power Plants," proceedings of the 35th Annual Western Protective Relay Conference, Spokane, WA, October 2008.

[2] E. O. Schweitzer, III, D. Whitehead, A. Guzmán, Y. Gong, and M. Donolo, "Advanced Real-Time Synchrophasor Applications," proceedings of the 35th Annual Western Protective Relay Conference, Spokane, WA, October 2008.

[3] D. Miller, R. Schloss, S. Manson, S. Raghupathula, and T. Maier, "PacifiCorp's Jim Bridger RAS: A Dual Triple Modular Redundant Case Study," proceedings of the 11th Annual Western Power Delivery Automation Conference, Spokane, WA, April 2009.

[4] W. Allen and T. Lee, "Flexible High-Speed Load Shedding Using a Crosspoint Switch," proceedings of the 32nd Annual Western Protective Relay Conference, Spokane, WA, October 2005.

[5] S. Shah and S. Manson, "Automated Power Management Systems for Power Consumers With On-Site Generation," presented at the 16th Annual Joint ISA POWID/EPRI Controls and Instrumentation Conference, June 2006.

X. VITAE

E. Roy Hamilton received his undergraduate degree in Electrical Engineering from the Georgia Institute of Technology in 1980. He has worked as an electrical design engineer for the Tennessee Valley Authority and as an I&E project engineer for E. I. DuPont and the former Amoco Oil Company. He also served as an electrical operations engineer for nine years in Saudi Aramco's Northern Area Producing Department. Currently, he is working in Chevron's Engineering Technology Center as an electrical engineering power systems specialist. He is a member of the IEEE and IEEE IAS.

John Undrill (F '78) received B.E. and Ph.D. degrees from the University of Canterbury, New Zealand. He joined General Electric in 1966 and since then has worked with General Electric, Power Technologies Incorporated, and Electric Power Consultants Incorporated. He has been fully engaged with power system dynamics issues throughout his career. He retired from corporate involvement in 2006 but continues to be active as a consultant.

Paul S. Hamer (S '70, M '74, SM '89, F '97) received the B.S.E.E. degree from Virginia Polytechnic Institute and State University, Blacksburg, and the M.S.E.E. degree from Oregon State University, Corvallis, in 1972 and 1979, respectively. In 1979, he joined Chevron Corporation, where he is currently a Senior Consulting Engineer, Electrical Machinery and Power Systems, with Chevron Energy Technology Company, Richmond, CA. His primary responsibilities include power system, motor, and generator application and consultation. He has worked on many refining, chemical, and oil production projects during his career with Chevron. In 2008, he was appointed a Chevron Fellow, Chevron's highest recognition for individual accomplishment. He has contributed to the American Petroleum Institute (API) standards for induction and synchronous machines and the API recommended practice on electrical area classification. He represents the API on the National Electrical Code, Code-Making Panel 11, on the subjects of motors, motor circuits, and controllers, and on the technical committee for NFPA 70E, Standard for Electrical Safety in the Workplace. From 1972 through 1977, he was with Westinghouse Electric Corporation, where he was a Service Performance Engineer with the Large Generator Department and an Industrial Power System Engineer and Resident Engineer with the Industry Services Division.

Scott Manson is a supervising engineer for the Engineering Services Division of Schweitzer Engineering Laboratories, Inc. (SEL). He received a Masters in Electrical Engineering from the University of Wisconsin–Madison and his Bachelors in Electrical Engineering from Washington State University. Scott worked at 3M Corporation as a control system engineer for six years prior to joining SEL in 2002. Scott has experience in designing and implementing control systems for electrical utility customers, high-speed web lines, multiaxis motion control systems, and precision machine tools. Scott is a registered professional engineer in Washington, Alaska, North Dakota, and Louisiana.