LBNL-2001105

Primary Frequency Response and Control of Power System Frequency

Author:

John Undrill

John Undrill, LLC

Energy Analysis and Environmental Impacts Division

Lawrence Berkeley National Laboratory

February 2018

This work was supported by the Federal Energy Regulatory Commission, Office of Electric Reliability, under interagency Agreement #FERC-16-I-0105, and in accordance with the terms of Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy.

Primary Frequency Response and Control of Power System Frequency

Prepared for the
Office of Electric Reliability
Federal Energy Regulatory Commission

Principal Author
John Undrill, John Undrill, LLC

Ernest Orlando Lawrence Berkeley National Laboratory
1 Cyclotron Road, MS 90R4000
Berkeley CA 94720-8136

LBNL-2001105

February 2018

The work described in this study was funded by Federal Energy Regulatory Commission, Office of Electric Reliability, under interagency Agreement #FERC-16-I-0105, and in accordance with the terms of Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy.

Acknowledgements

The work described in this study was funded by the Federal Energy Regulatory Commission (FERC), Office of Electric Reliability, under interagency Agreement #FERC-16-I-0105, and in accordance with the terms of Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The authors acknowledge project management provided by Eddy Lim, FERC Office of Electric Reliability. The author thanks Sydney Niemeyer, NRG Energy (retired), and Julia Matevosyan, Electric Reliability Council of Texas, for review comments on an early draft of this report.

All opinions, errors, and omissions remain the responsibility of the authors. All reference URLs were accurate as of February 2018.

Table of Contents

Table of Figures

List of Tables

Acronyms and Abbreviations

AGC	Automatic generation control
CAISO	California Independent System Operator
DCS	Digital Control System
DOE	Department of Energy
Efrac	Electronically-coupled fraction of generation
ERCOT	Electric Reliability Council of Texas
ERSWG	Essential Reliability Services Working Group
FERC	Federal Energy Regulatory Commission
GE	General Electric
GW	Gigawatt
Hz	Hertz
LR	Load Resource (load shedding used as PFR)
LBNL	Lawrence Berkeley National Laboratory
mHz	Millihertz
NERC	North American Electric Reliability Corporation
Nfrac	Non-responsive fraction
NREL	National Renewable Energy Laboratory
PFR	Primary Frequency Response
PLC	programmed logic controller
Rfrac	Responsive fraction
ROCOF	rate of change of frequency
Sfrac	Responsive-sustaining fraction
UFLS	Under Frequency Load Shedding (emergency load shedding)
WECC	Western Electricity Coordinating Council

Terminology and Definitions

Balancing Control	The action of controllers in the greater power system that supervise the plant secondary controllers. These controls typically act by transmitting reference setpoints to plant secondary controllers
Base MVA	The base value of generator MVA used in stating variables describing an machine in per unit form with respect to the generator
Base Power	The base value of prime mover (turbine) power used in stating variables describing a machine in per unit form with respect to the prime mover
Electronic fraction (Efrac)	Fraction of generation that is electronically coupled and does not respond to change of frequency
H, Inertia Constant	The inertia constant of an individual rotating machine stated in seconds with respect to generator base MVA
Non-responsive fraction (Nfrac)	Fraction of generation that is rotating, synchronous, contributes inertia, but does not respond to change of frequency
Primary Control	The action of the governor in response to change of detected speed and/or power
Responsive fraction (Rfrac)	Fraction of generation that is responsive to change of frequency in accordance with governing droop (see Figure 4).
Responsive-sustaining fraction (Sfrac)	Fraction of the responsive generation that sustains its initial response in accordance with the governor droop
Secondary Control	The action of a controller at a generating plant that supervises the governor and acts by changing the governor speed-load reference
System Inertia	The summation, $\Sigma(H * Mbase)$ of all rotating machines connected to the synchronous system

Note: The fractions, Efrac, Nfrac, Rfrac, and Sfrac are illustrated in Figure 9.

1. Introduction

This report is a review of factors that affect the ability of a large synchronous electric power system to control and manage its frequency. This report is an extension of Undrill 2010[1]; it uses the same overall approach to power system dynamics as used in Undrill 2010, but uses greater detail in its description of the power plants making up the generation fleet and, based on the greater detail, addresses factors that Undrill 2010 touched on only in principle.

2. Stipulations

This report is concerned with the control of generating plants with respect to the frequency of the power system. It recognizes that the dispatch of power in an interconnected power system takes account of electric transmission limitations, hydraulic limitations of river systems, availability of reactive power support, and many other issues. For the consideration presented here, however, it is stipulated that the electric transmission system connecting the fleet of power plants is adequate for the task of holding all generation in synchronism and that generation fleet is dispatched with proper regard for transmission system limitations. Accordingly, the transmission system is not modeled in the simulations used in this study. Rather the transmission grid is treated as a single electrical bus with the whole generation fleet connected to it as indicated by Figure 1.

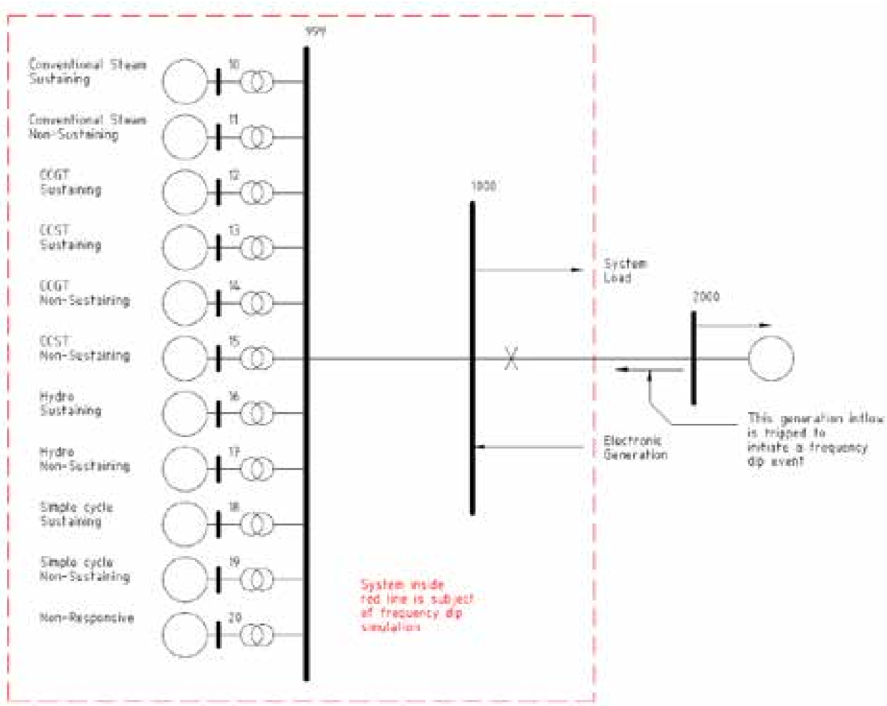

Figure 1. Generation connected at an electrical point

[1] Undrill, J.M. (2010). *Power and Frequency Control as it Relates to Wind-Powered Generation*. LBNL-4143E. https://certs.lbl.gov/publications/power-and-frequency-control-it

3. The Basic Power System Operation

3.1 Balancing Control

It is desired to operate the electric power system with its frequency as nearly constant as practical and with the real power outputs of individual power plants in accordance with a desired dispatching process. In the most basic form the dispatching process allocates power output targets to generators with recognition of overall system operating cost and of operational limitations as varied as transmission line capacities, river flow requirements, and air quality.

The focus of this report is the control of frequency with the stipulation that electric transmission and other limitations, while present, are a secondary consideration. Before putting such issues aside, however, we should consider the following simplified but illustrative example.

Consider three generators, each connected to a local load and also connected to the other two generators as shown in Figure 2. The costs of producing power at the three generators are not equal and the 'invisible hand of the market' dictates that the individual generators should be operated at outputs that differ from the adjacent loads. Consider that Generator 1 can produce power at an attractively low price while power from Generator 3 is expensive. This will lead to power flows as shown in the figure. It is clear that the supply of power to the load at Site 3 can take advantage of the lower cost power offered at Site 1 only up to the flow limit on the transmission.

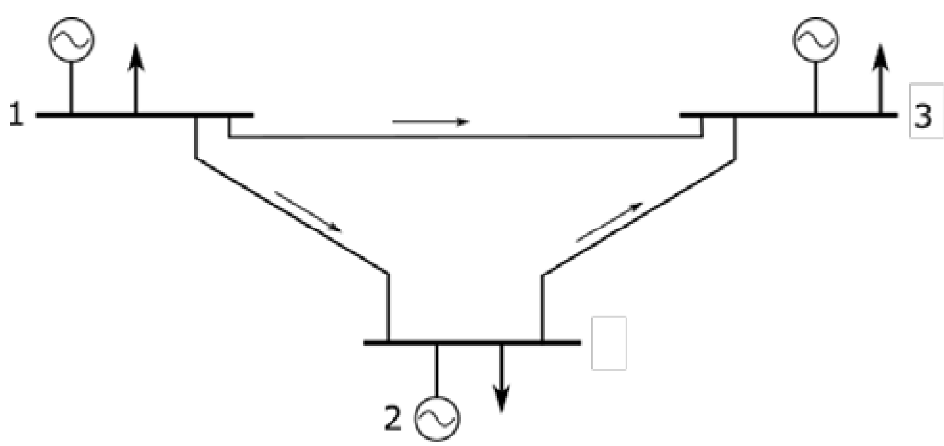

Figure 2. Power flows across an electric transmission system

Generator outputs are adjusted minute-by-minute to balance total power production with total load. The calculations of the desired generator outputs are done in central dispatch offices and updated real power output targets are transmitted to the generators every few seconds.

The process of adjusting generator power output targets is referred to throughout this report as balancing control. While it is in essence the control that could be achieved by an attentive human operator watching a host of indicating instruments and 'pulling levers' with expertize, the balancing control undertaking in a large grid is now of such scale and complexity that it is implemented by digital computers working on a time scale of tens of seconds to tens of minutes.

This report recognizes that balancing controls are active and that the power output setpoints that they provide to power plants are updated from time to time. With this recognition, however, the discussion of system dynamic behavior in this report stipulates that the targets allocated to generating plants by system dispatch processes are constant for the periods under consideration.

3.2 Power Plant Control

With the current state of the power dispatching process, in combination with present load forecasting techniques, balancing control action is able to keep the total real power production and load in remarkably good, but not perfect balance. There is always some imperfection in the dispatch process.

Further, the balance between power production and consumption is exposed to disturbances that are foreseeable but whose times of occurrence cannot be predicted.

The frequency of the power system, and the flows on its transmission lines, respond to variations of load on a time scale of tenths of a second to seconds. To maintain control of system frequency, the power outputs of the prime movers in generating plants (turbines, not generators) must be changed on this time scale.

The response of generators to suddenly occurring events, such as the tripping of a generator, is much too quick to be managed by balancing controls. Accordingly, the principle of power system control is that:

- balancing controls instruct power plants as to what they are to do, but not how to do it
- the doing of what balancing control instructs is assigned to primary and secondary controls within the individual power plants
- the quick response that is essential for the maintenance of stable control is handled by primary controls

In most rotating power plants the primary controls of interest are the governors which are tightly linked to the turbines, and the secondary controls are the supervisory computer systems that run the operator consoles and the plant systems other than the turbines.

3.3 Power and Frequency Control

Consider a power system consisting of two turbine generators connected by a transmission line as shown in Figure 3. The loads at the two ends of the system are variable and the transmission line has limited capacity. It is necessary to run the system at its nominal speed (or frequency) and to maneuver the turbines to keep the transmission line flow within its limits. The two ends of the system are many miles apart; communication between the ends is slow and less than perfectly secure.

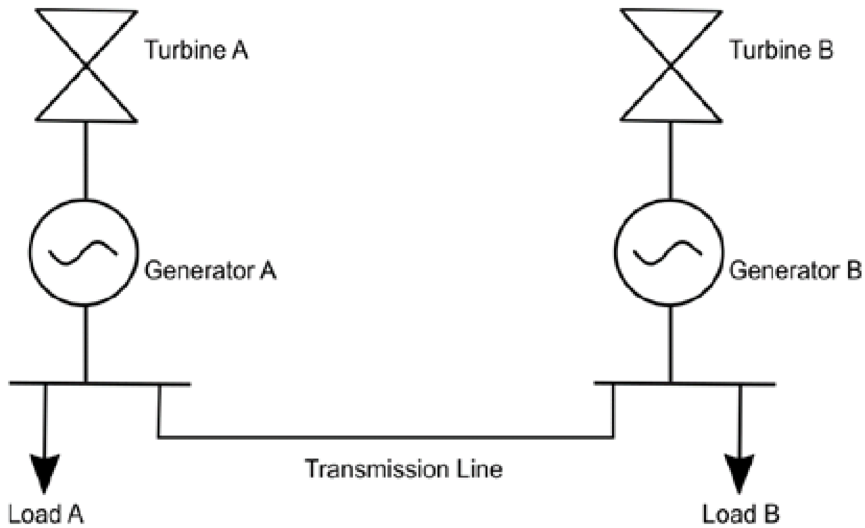

Figure 3. Synchronous power system with two turbine-generators

The turbine-generator at each end of the system is controlled by a governor whose function is to make its speed and power lie on a sloping line in the graph of speed versus power, as shown in Figure 4.

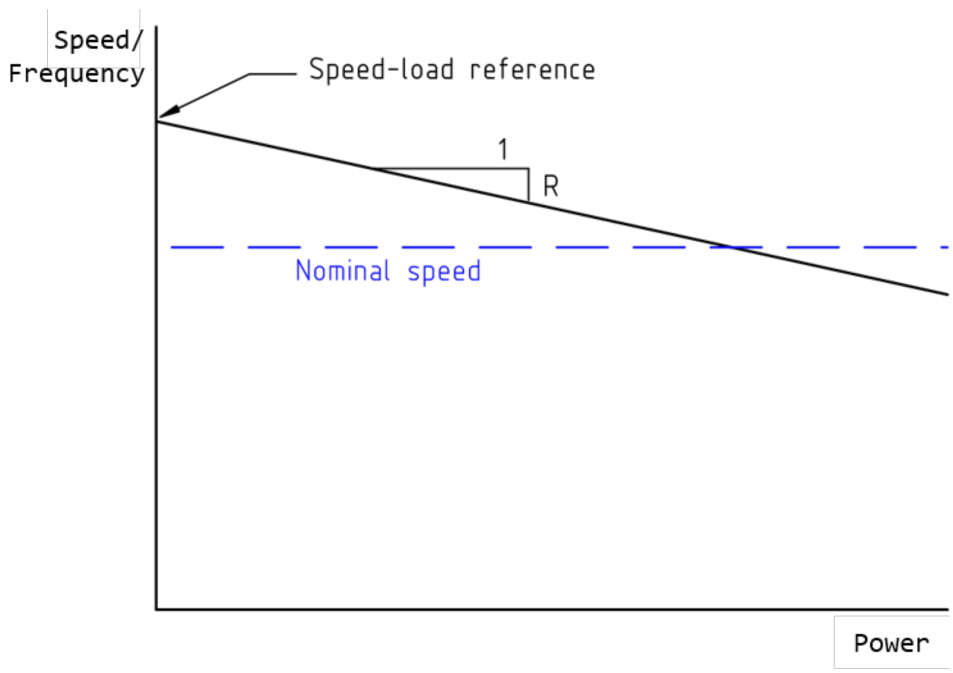

Figure 4. The governor–droop relationship

The system is operated as illustrated by Figure 5 through Figure 7. First consider that the system is running at nominal speed with the two turbines producing the powers *OA* and *OB* as shown in Figure 5. The two governors have a common droop value, *R*, and speed-load reference settings are as indicated by *a* and *b*. The total generation and total load are equal, and equal to (*OA*+*OB*).

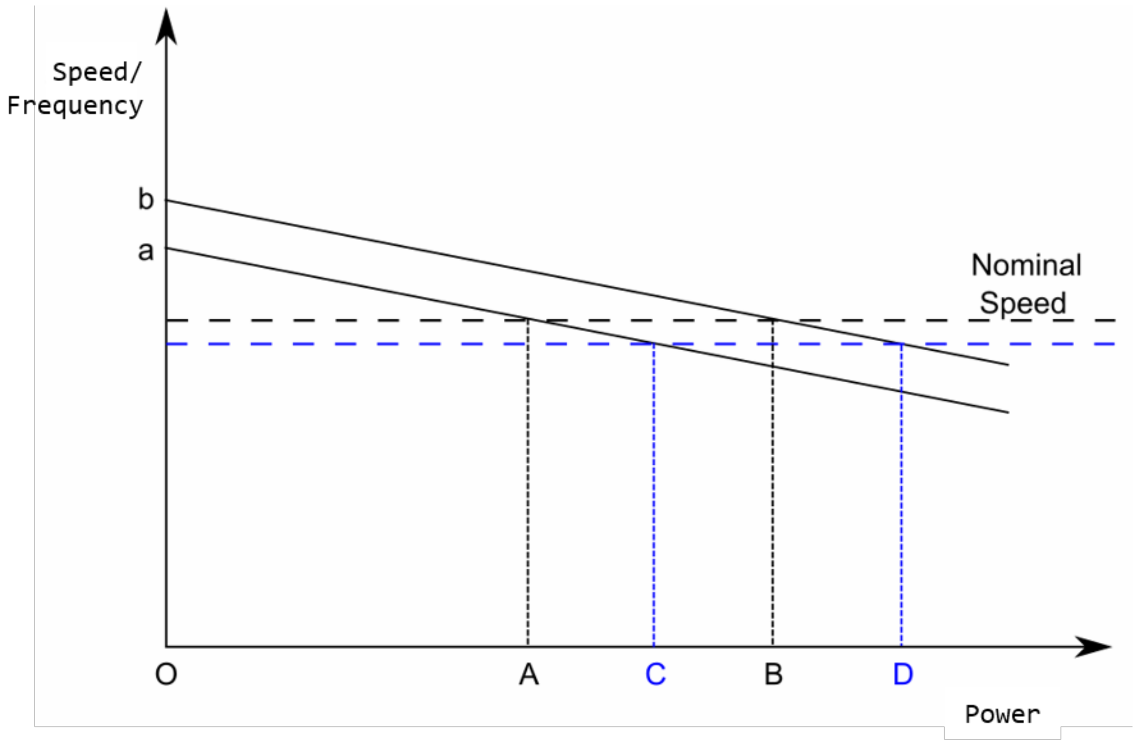

Figure 5. Power system control; Governor response to increase in system load

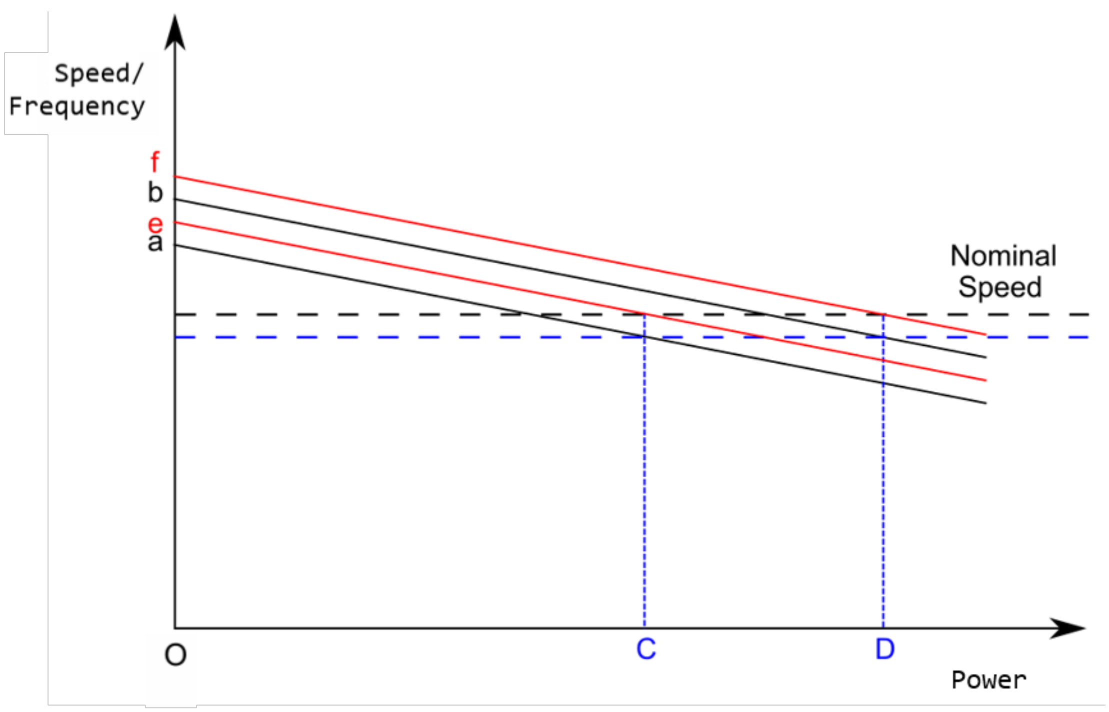

Figure 6. Power system control; Increase speed-load reference settings equally *a* to *e* and *b* to *f* to raise frequency while keeping total power and individual power output unchanged

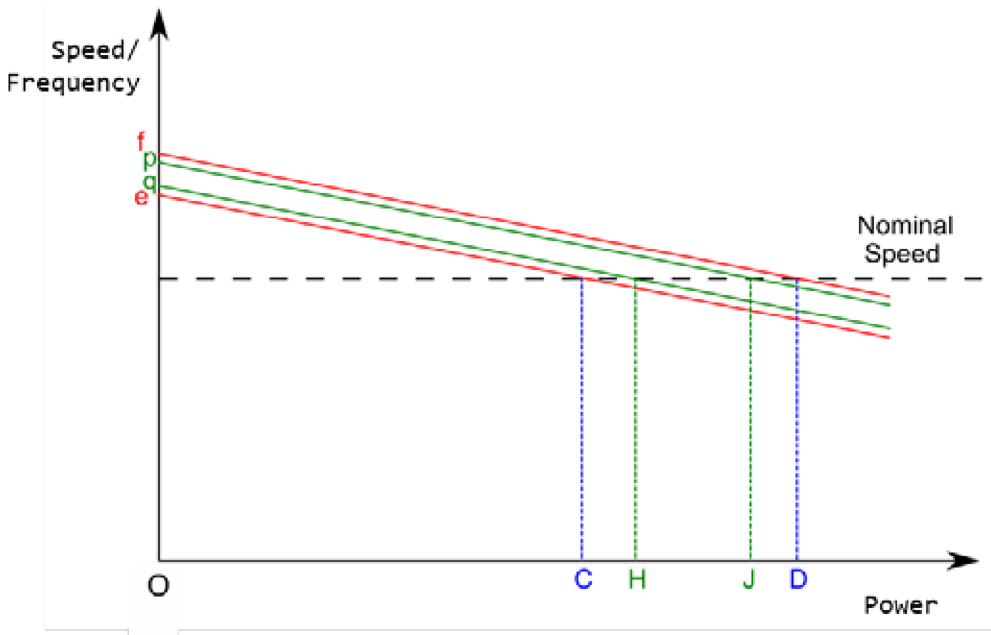

Figure 7. Power system control; Increase speed-load reference *e* to *q* and decrease speed-load reference *f* to *p* to increase power of Turbine A and reduce power of Turbine B

Now consider that the total load increases, without it being known where the increase occurs. The total generation must increase to match the new total load. The governor speed-load reference settings, *a* and *b*, are not changed. The frequency must decrease so that the generator outputs increase to the new values *OC* and *OD*. The realignment of frequency and turbine powers occurs in a few seconds and operation continues. The balance of generation and load is correct, but the frequency is slightly below the desired nominal value.

Frequency can be restored by adjusting the two governor speed-load references as shown in Figure 6. The reference settings are increased equally to *e* and *f* so that the frequency is again at nominal value.

Now consider that it is necessary to increase the output of Generator A and, correspondingly, reduce the output of Generator B. Figure 7 illustrates how this is done. The speed-load reference, *q*, of Turbine A is increased and the speed-load reference, *p*, of Turbine B is lowered. The turbine powers settle at the new values *OH* and *OJ*, such that the total power (*OH*+*OJ*) = (*OC*+*OD*) is unchanged.

These three figures are a complete description of how a power system is operated by its primary, secondary, and balancing controls.

Governors are primary control. They do not control speed and they do not control power; rather, they control the relationship between speed and power. The relationship is depicted by the sloping lines in the figures and is given formally by

$$\omega = \omega_{ref} - RP$$

Equation 1

where

ω is the running speed of the turbine
ω_{ref} is the speed-load reference setting of the turbine governor
P is the power output of the turbine (or, as convenient, of the generator that it drives)
R is the governor droop setting

With the turbine governors throughout the system operating to an agreed droop setting (typically R=0.04 or R=0.05) the system is controlled by adjusting the governor speed-load references.

- Collective increases and decreases of speed-load references act to increase or decrease system frequency
- Differential adjustments of speed-load references act to change the distribution of power among the turbine-generators

This method of control is long established. It can operate very successfully with no automated central controller, does not depend on broad bandwidth communication between power plants, and is scalable from the smallest power systems to the largest.

3.4 Summary

- Turbine governors are primary control and respond immediately to changes in frequency and power. Turbine governors operate autonomously but, by agreement, accept instructions from plant secondary controls and system balancing controls.
- Power plant internal control systems (digital control systems, or DCS) are secondary control. They instruct the turbines in accordance with the requirements of the plant and instructions received from system balancing controls. The instructions are delivered as adjustments of the governor speed-load references.
- Balancing controls are the management and optimization system by which the entity responsible for the interconnected generating and transmission system maintains order among widespread and disparate entities. The instructions of balancing controls are delivered to power plants by a wide variety of automatic and manual means.

Primary controls, the turbine governors, have absolute authority within their sphere of control and on their time scale. Secondary controls have absolute authority within their spheres and on a longer time scale than that of primary controls. Balancing controls depend on primary and secondary controls for the stability and continuity of the system but have authority over the dispatch of generation to meet the requirements of the electric transmission system and commerce.

The remainder of this report is concerned with the ability of primary controls, turbine governors, to maintain prompt and stable control of the power system in normal and emergency conditions.

4. Primary and Secondary Controls

The control of a large power system is hierarchical. In this discussion we consider the actions of *primary control* and *power plant secondary control*. We recognize the existence of *grid-level balancing controls* and *grid-level tertiary controls*; we recognize that the actions of grid level controls are the means by which the power system is maneuvered to and from the conditions that we consider, but our focus is on primary and power plant secondary control action.

4.1 Primary Control

In relation to grid frequency, primary control is the immediate control of the relationship between turbine speed[2] and power. In rotating machines it is exerted by governors that manipulate control valves. In electronically coupled generation it is exerted within the electronics. (Wind turbine-generators may have primary control exerted by either of both of governors and electronic controllers. The details of wind turbine primary control are not considered here; wind turbines are considered to be non-responsive electronically coupled generation for the discussion presented in this report.) Primary Control Response is the change of power output produced by the governor or electronic controller in response to a change in turbine speed or grid frequency. The grid is operated on the basis that the primary controls of power plants are autonomous and will function continuously to enforce the required relationship between speed and power without requiring inputs from any external source. This form of primary control is depended on, millisecond-by-millisecond, as the fundamental phenomenon that keeps the grid in stable equilibrium.

4.2 Power Plant Secondary Control

Secondary control is local to a power plant; it is the control by which turbines respond to commands originating in the plant. Secondary control inputs are normally applied to the governor speed-load references. It is common for a modern power plant to have several distinct modes of secondary control and, also, to be able to accept inputs from external sources such as the Automatic Generating Control[3] (AGC) system of a Balancing Area. For example, a multi-unit hydro plant would typically have a plant controller whose objective is to optimize the allocation of loading among the several turbines to make best use of water. This plant controller would be able to accept a setpoint for total plant output either as a local manual operator input or from the AGC system at a Balancing Area control center (via telecontrol). In addition, there is the possibility of the telecontrol from the grid control center delivering input signals directly to the governor speed-load references, bypassing the plant controller. All of the many ways of manipulating the governor references are secondary control in that they are means by which plant and grid operators can control the turbines by the issuance of instructions.

[2] In electronically controlled plants where there is no mechanical prime mover, the frequency of the electronic converter takes on the role of 'speed'.

[3] Automatic Generating Control (AGC) is a form of balancing control.

4.3 Grid-Level Balancing Control

Power plant secondary controls are supervised by balancing control action at the grid level. Grid level controls are typically implemented at transmission system dispatch offices and involve levels of automation ranging from continuous-acting feedback loops to human observation and verbal instructions. The AGC systems that are used to regulate real power flows between Balancing Areas are grid-level balancing controls.

5. Power Plant Operating Modes

The control service that a generating plant makes available to the grid, for both primary and secondary control, is determined by the ways the individual power plant operator(s) choose to run their machines. Few turbines are controlled manually, though manual control is possible in some special situations. The great majority of the turbines on the grid are operated by unit and plant level control systems that can be put into a very wide variety of modes. The plant control modes that account for the majority of the generating capacity and are of interest here are:

a. *Simple droop mode.* The governor receives secondary control inputs only by manual actions of the turbine operator. A turbine running in simple droop mode provides primary control response to the grid but no automatic or dependable secondary response. In the terminology set out below turbines operating in this mode are responsive-sustaining; they are in the part of the rotating generation fleet designated by the fraction *Rfrac*.

b. *Prescheduled output mode without frequency bias.* A plant controller applies secondary control commands to the governor speed-load reference to hold the plant at a prescheduled output without reference to grid frequency. This prescheduled output typically is a constant or a ramp at a preset rate. In this mode a turbine provides primary response to a change of grid frequency changes temporarily, but is returned to its prescheduled output by the secondary control. In the terminology set out below turbines operating in this mode are responsive/non-sustaining; they are in the part of the rotating generation fleet described by the fractions *Rfrac, and 1-Sfrac*.

c. *Prescheduled output mode with frequency bias.* A controller applies secondary control commands to the governor speed-load reference to hold the plant at a prescheduled output with the prescheduled output being biased by deviation of grid frequency. The prescheduled output is stated as the output to be produced when grid frequency is at scheduled value. This prescheduled output typically is a constant or a ramp at a preset rate. In this mode a change in output made by primary control action is sustained as long as the frequency remains different from nominal value. In the terminology set out below turbines operating in this mode are *Responsive-sustaining*; with properly chosen frequency bias factors they are in the part of the rotating generation fleet described by the fractions *Rfrac* and *Sfrac.*

d. *Automatic generation control mode.* (Also referred to as Load-Frequency control mode) The scheduled power output of the turbine is manipulated by signals received from the AGC system of the Balancing Area. In this mode the turbine provides primary response and whether this primary control response is sustained or not depends on the action of the AGC System. Turbines receiving speed-load reference signals directly from grid-level balancing controls are considered here to be in the fraction of the rotating fleet described by *Rfrac*.

e. *Non-responsive mode.* The turbine control valves are wide open, are at a fixed position, or are under the command of a controller that does not respond to turbine speed or grid frequency, such as the exhaust temperature limiter of a gas turbine or the pressure controller of a steam turbine. For the purpose of this discussion a non-responsive turbine provides neither primary nor secondary control response. Non responsive turbines are in the part of the rotating fleet described by the fraction *Nfrac*.

A key distinguishing difference between primary and secondary control is that a power system would be able to operate steadily, for a few minutes at least, with no automatic secondary control action or balancing control action being taken, but that it would not be possible for it to operate at all without the continuous exertion of effective primary control.

Another useful distinction is that primary control is the immediate feedback control function that regulates the power of the turbine on a time scale that is much quicker that a human operator could achieve, while secondary control is most often the automation by a computer of a function that could be done by an attentive human operator.

A last distinction is that primary control is the enforcement of a simple and purely local control objective for each turbine individually, while secondary control is concerned with managing the relationships between multiple turbines. Both levels of control affect transmission flows; this effect is a byproduct of primary control action and one of the main objectives of grid-level secondary control.

The power system requires both Primary Control Response and Secondary Control Response from turbines connected to it. It is neither possible nor necessary for all turbines to contribute to these control responses, but a sufficient quantity and appropriate geographic distribution of each is essential. Primary control is the essential means by which the grid is assured to be stable and controllable. Secondary control is used to manage resources. Secondary control systems are used to maneuver power plants in accordance with a broad field of considerations including energy markets, transmission security requirements, and internal operational necessities of individual power plants. For this discussion, the overriding interest in secondary control is its role in ensuring that the power system can match its generation to its load and can maintain control of its frequency.

6. Power System Model

6.1 Generation

The generation is divided as follows:

- *Responsive synchronous generation*. Generation operating in one of the modes (a) through (d) noted in Section 5.
- *Non-responsive synchronous generation*. Generation operating in mode (e) noted in Section 5.
- *Electronically coupled generation*. Generation whose output is not sensitive to frequency or electrical phase angle, and which has no rotating elements that contribute to system total inertia.

The responsive synchronous generation is subdivided as follows into two categories depending on the operating modes in effect in the power plants:

- *Responsive-sustaining generation*. Synchronous rotating generation operating in simple droop mode that is able and willing to sustain changes made in response to a frequency change for an extended period (typically many minutes).
- *Responsive/non-sustaining generation*. Synchronous rotating generation operating in a mode such that it responds to a change of frequency in substantially with same way as responsive-sustaining generation but whose change of power output is not sustained, even though the frequency change persists.

6.2 Makeup of the Generation Fleet

6.2.1 Participation in frequency control

The simulation model is made up on a per unit basis. The base of the per-unit values is MW_{init}. MW_{init} is the total generation of the system *at the moment, t^-, just before a block of generator power has been interrupted but excluding the generation that is interrupted*.

- The responsive synchronous generation produces the fraction, *Rfrac*, of the total power generation.
- The fraction, *Sfrac* of the responsive synchronous generation is responsive-sustaining generation.
- The fraction (1–*Sfrac*) of the responsive synchronous generation is responsive/non-sustaining generation.
- The electronically coupled generation produces the fraction, *Efrac*, of the total power generation.
- The balance, *Nfrac*, of the generation required to meet the system load is produced by non-responsive synchronous generation.

- The fleet makeup fractions are such that

$$Rfrac + Efrac + Nfrac = 1$$

Equation 2

$$Sfrac < 1$$

Equation 3

For the simulations shown in this report, the responsive fraction was assigned a common value across all types of generation.

6.2.2 Reserves for control action

The simulations presented in this report were made with all prime movers in the responsive part of the system loaded initially at 91 percent of their maximum outputs. Given this:

- The maximum response to a change of frequency in the short term, while the contribution of responsive-non-sustaining generation is present, is equal to $0.09 * Rfrac * MW_{init}$ megawatts.
- The maximum response to a sustained deviation of frequency in the long term, after the contribution of responsive-non-sustaining generation has been withdrawn, is equal to $0.09 * Sfrac * Rfrac * MW_{init}$ megawatts.

6.2.3 Loss of generation disturbances

The fleet, just prior to the loss of generation, has total load and total generation output of $(1 + \Delta P)$ times the total output of the generation fleet described by *Rfrac*, *Sfrac*, *Efrac* and *Nfrac*. The magnitude of the generation loss, ΔP, is therefore specified as a fraction of the total generation produced immediately before the event by the generators that are not tripped.

6.3 Makeup of the Generation Fleet: Plant Characteristics

The rotating synchronous generation fleet is modeled by eleven blocks of generation as listed in Table 1. The different types of generation respond differently to a transient variation of frequency. Figure 8 shows the simulated response of a large steam turbine, a large industrial gas turbine, and a hydro turbine, to a representative transient dip in grid frequency.

Two alternative fleet profiles are considered; a predominantly thermal fleet whose makeup by plant type is given in Table 2 and a mixed thermal-hydro fleet described by Table 3.

The response characteristics of the difference blocks of generation are modeled by the dynamic simulation models listed in Table 4.

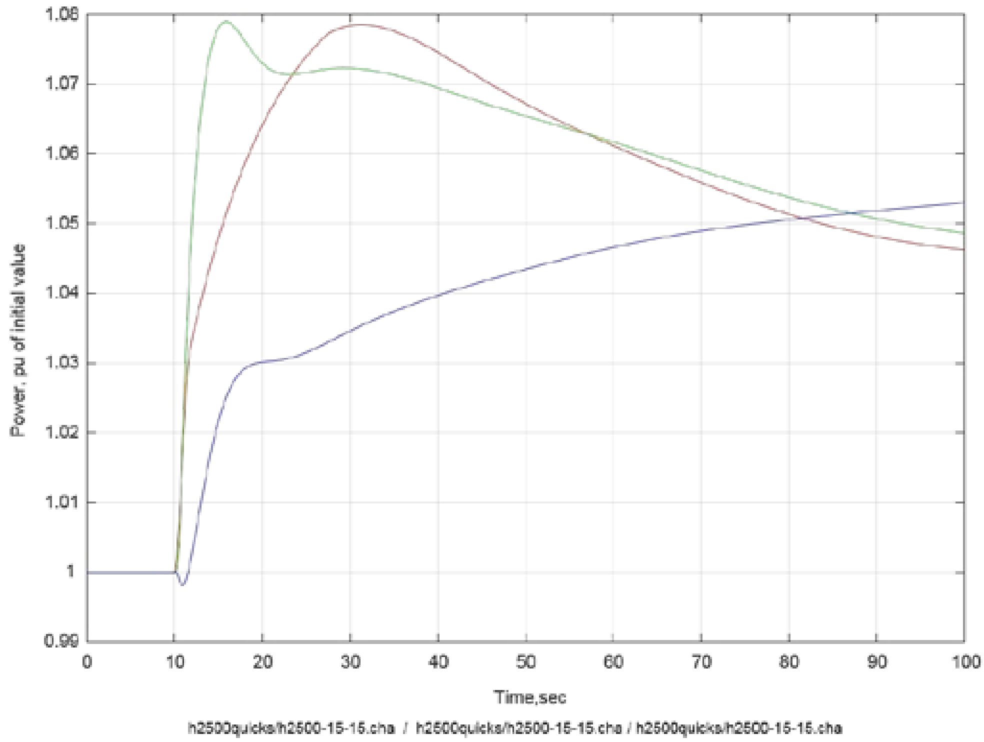

Figure 8. Comparison of thermal plant and hydro plant primary response to loss of 0.02 per unit generation

Red - 10pm1 Steam turbine response
Green - 12pm1 - Gas turbine response
Blue - 16pm1 - Hydro plant response

Table 1. Generation classes

Block	Name	Plant Type	Operating Mode	Note
10	SteamA	steam turbine	Responsive-sustaining	
11	SteamB	steam turbine	Responsive/non-sustaining	
12	CCGTA	combined cycle - GT	Responsive-sustaining	
13	CCSTA	combined cycle – ST	Responsive-sustaining	Slow response following GT
14	CCGTB	combined cycle – GT	Responsive/non-sustaining	
15	CCSTB	combined cycle - ST	Responsive/non-sustaining	Slow response following GT
16	HydroA	hydro	Responsive-sustaining	
17	HydroB	hydro	Responsive/non-sustaining	
18	AeroA	simple cycle GT	Responsive-sustaining	
19	AeroB	simple cycle GT	Responsive/non-sustaining	
20	NonResp	synchronous	Nonresponsive	
21	Elec	electronic	Electronically coupled generation	

Table 2. Make up of thermal system

Block	Name	Plant Type	Fraction, $Rfrac_{type}$, of Responsive Generation
10,11	Steam	steam turbine	0.40
12,14	CCGT	combined cycle - GT	0.36
13,15	CCST	combined cycle - ST	0.18
16,17	Hydro	hydro	0.04
18,19	Aero	simple cycle GT	0.02

Table 3. Make up of mixed thermal-hydro system

Block	Name	Plant Type	Fraction, $Rfrac_{type}$, of Responsive Generation
10,11	Steam	steam turbine	0.14
12,14	CCGT	combined cycle - GT	0.24
13,15	CCST	combined cycle - ST	0.12
16,17	Hydro	hydro	0.44
18,19	Aero	simple cycle GT	0.06

6.4 Modeling of Rotating Prime Movers

The rotating synchronous generation is represented in simulations by the dynamic models listed in Table 4. The lcfb1 model is turned off for turbines operating in responsive-sustaining mode and turned on for turbines in responsive/non-sustaining model.

Table 4. Prime mover and plant control modeling

Block	Name	Plant Type	Turbine Governor	Plant Load Controller	Sustains Response
10	SteamA	steam turbine	ggov1	lcfb1	Y
11	SteamB	steam turbine	ggov1	lcfb1	-
12	CCGTA	combined cycle - GT	ggov1	lcfb1	Y
13	CCSTA	combined cycle - ST	ccst3		Y
14	CCGTA	combined cycle - GT	ggov1	lcfb1	-
14	CCSTB	combined cycle - ST	ccst3		-
16	HydroA	hydro	h6e	lcfb1	Y
17	HydroB	hydro	h6e	lcfb1	-
18	AeroA	simple cycle GT	ggov1	lcfb1	Y
19	AeroB	simple cycle GT	ggov1	lcfb1	-
20	NonResp	synchronous	n/a	n/a	
21	Elec	electronic	n/a	n/a	

Except where noted, the standard forms of these models as implemented in the General Electric (GE) PSLF computer program are used. The electric system and dynamic model data of the power system model are contained in the data files associated with this report: s10-1.sav and s10-1.dyd.

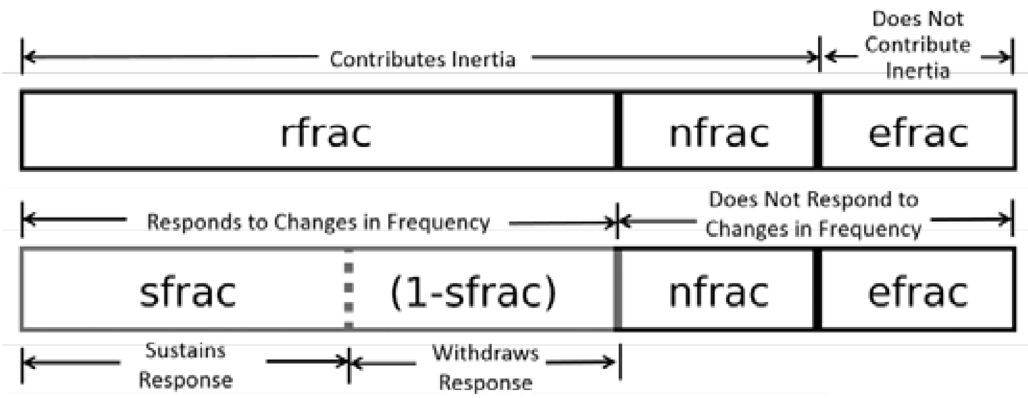

Figure 9. Makeup of the generation fleet

6.5 Load

Except where noted, the real power load is represented as a function of frequency as follows:

$$P_{load}(f) = P_{load}(f_{nom}) * (1 + D_l\,(f - f_{nom})$$

Equation 4

6.6 Load Shedding

6.6.1 Protective load shedding

The system model used for simulations shown in this report includes a representation of protective underfrequency load shedding. This is included in the model solely as a measure to allow simulations to continue to function in situations where frequency response of generation has failed to arrest a decline of frequency or has failed to sustain the arrest. Three stages of shedding are set to operate and disconnect one half percent of total system load, each, at 59.1, 59.0 and 58.9Hz.

Table 5. Protective load shedding settings

Underfrequency settings	59.1	59.0	58.9	Hz
Relay definite time delay	0.025	0.025	0.025	second
Circuit breaker delay	0.05	0.05	0.05	second
Load disconnected	0.005	0.005	0.005	per unit of initial total load

6.6.2 Load shedding used for primary frequency response

Simulations shown in Section 13 consider the use of preprogrammed load shedding as a means of obtaining frequency response at frequencies well above the level at which protective load shedding comes into play. This use of load shedding, such as the LR (Load Resource) provisions of the ERCOT system, is modeled by underfrequency relays with frequency and delay settings as noted in Section 13.

7. Factors Affecting the Control of Grid Frequency

7.1 Initiating Loss of Generation

Excursions of power system frequency from the desired value are the result of changes of load or generation that take place too quickly to be handled by the systems' balancing controls. The changes of greatest interest are the sudden loss of a large block of generation, the sudden connection of a large load, and the sudden loss of a large block of load. All can be studied, for the purpose of understanding the dynamic behavior of the system, by simulations of the inertial behavior of the rotating equipment of the system and the associated behavior of primary controls.

The initiating event in all simulations used for this report is an instantaneous loss of generation. The size of the loss is either two percent, or four percent, of the collective output of the generating fleet (rotating plus electronic) immediately after the loss. In relation to the total generating capacities of the three major U.S. interconnections these loss of generation events are as shown in Table 6.

Table 6. Generation loss event sizes in MW for the major synchronous systems

System	Total (MW*1000)	2% Loss (MW)	4% Loss (MW)
WECC Max Load	120	2400	4800
WECC Min Load	60	1200	2400
ERCOT Max Load	60	1200	2400
ERCOT Min Load	30	600	1200
Eastern Max Load	400	8000	16000
Eastern Min Load	200	4000	8000

7.2 Summary of parameter-range simulation series

Table 7 and Table 8 summarize several series of simulation runs that are discussed below.

Table 7. Simulation series mapping frequency nadir as function of *rfrac* and *sfrac*

Name	Generation Loss %	Electronic Generation Fraction	Rotating Fleet Makeup	Figure	Unit Inertia Constant	System Inertia (MW/sec)
a1000r*	2.0	10	Thermal	Figure 22	4	43200
a2500r	2.0	25	Thermal	Figure 23	4	36000
a4000r	2.0	40	Thermal	Figure 24	4	28800
b0100r	4.0	1	Thermal	Figure 21	4	39600
b1000r	4.0	10	Thermal	Figure 21	4	43200
p2500s	2.0	25	Thermal	Figure 10	4	27000
h1000q‡	2.0	10	Mixed-hydro	Figure 25	3	32400
h2500q	2.0	25	Mixed-hydro	Figure 26	3	27000
h4000q	2.0	40	Mixed-hydro	Figure 27	3	21600

** Base Case*
‡ Quick reset on hydro, Ki=0.02

Table 8. p2500 series runs

File	Δ Per Unit	Efrac	Rfrac	Sfrac	Ki	H	Tact	Note
p	0.02 - 0.04	0 - 0.5	0.9 - 0.4	0.9	0.02	4	0.2	Vary responsive fraction
q2	0.02 - 0.04	0 - 0.5	0.4	0.9	0.02	4	0.2	As p2 runs but Rfrac = 0.4
q3	0.02 - 0.04	0 - 0.5	0.4	0.9	0.02	3	0.2	As q2 but reduced H from 4 to 3
r2	0.02	0 - 0.5	0.4	0.9	0.0067 - 0.02	4	0.2	Vary secondary response rate
r3	0.02	0 - 0.5	0.4	0.9	0.0067 - 0.02	3	0.2	As r2 runs but reduced H from 4 to 3
s2	0.02	0 - 0.5	0.4	0.6	0.0067 - 0.02	4	0.2	As r2 runs but Sfrac = 0.6
s3	0.02	0 - 0.5	0.4	0.6	0.0067 - 0.02	3	0.2	As s2 runs but reduced H from 4 to 3
t2	0.02 - 0.04	0 - 0.5	0.4	0.9	0.02	4	2.0	As q2 runs but slow governing response
u2	0.02 - 0.04	0 - 0.5	0.9 - 0.4	0.9	0.02	4	0.2	As q2 runs but reg margin = 0.04
v2	0.02 - 0.04	0 - 0.5	0.9 - 0.4	0.9	0.02	4	0.2	As q2 runs but reg margin = 0.06

8. Base Case Example

All simulations discussed in this section use the thermal base case system model described in Section 6.3. Ninety percent of the generation is rotating machines whose inertia constants are in the conventional range; the remaining ten percent is electronically coupled equipment that works at constant power. Forty percent of the total generation is responsive to change of frequency and ninety percent of that responsive generation sustains increases in output. (*Rfrac* = 0.4, *Sfrac* = 0.9)

The system loses two percent of its generation. Frequency starts to decline immediately. Governors respond, the decline of turbine speeds is arrested, and frequency is stabilized at an acceptable value. Figure 10 shows four different aspects of the system's response.

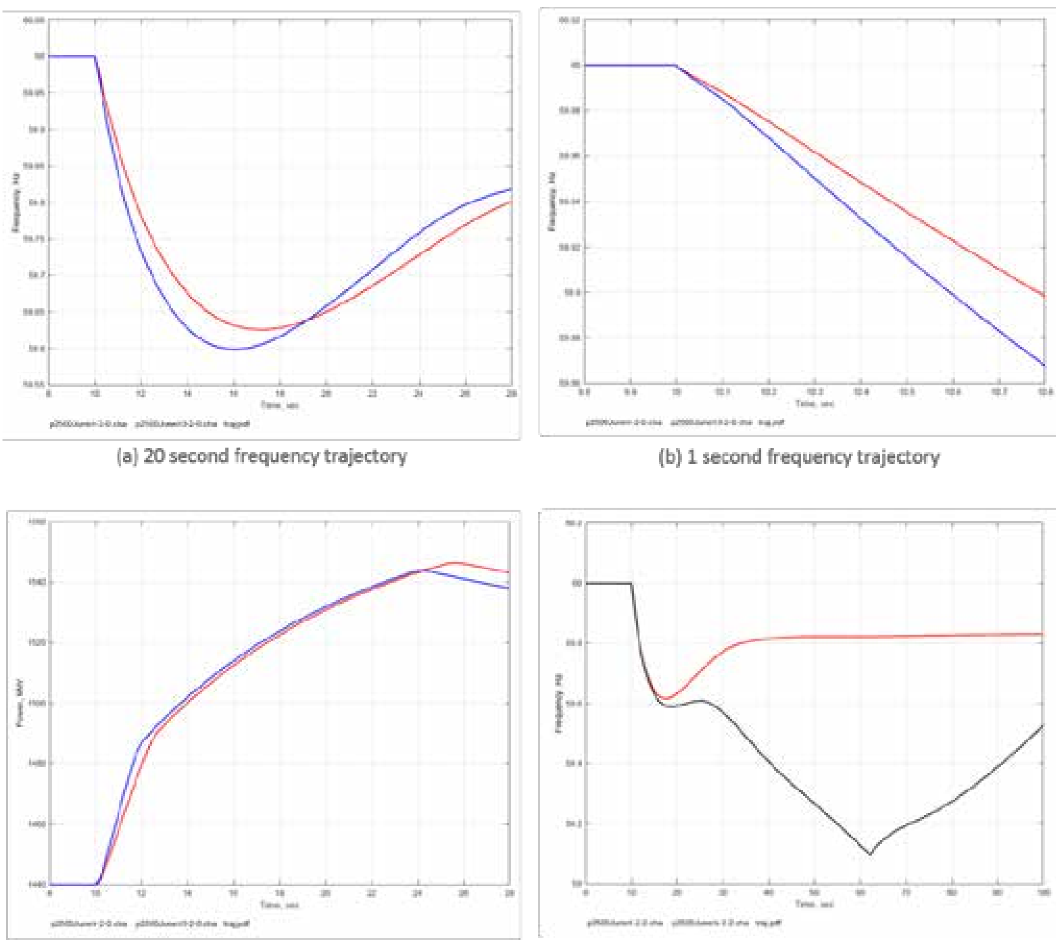

(a) 20 second frequency trajectory

(b) 1 second frequency trajectory

(c) 20 second power trajectory

(d) 100 second frequency trajectory

Figure 10. Four views of a frequency trajectory

The red curve in Figure 10(a) shows a commonly seen view of the trajectory of frequency versus time. The nadir of frequency is at 59.63Hz and after 20 seconds frequency has recovered to 59.8Hz. All appears to be well. However, Figure 10(a) is far from the complete description of the event and tells much less than the complete story.

The red curve in Figure 10(a) was made assuming that all of the rotating generators have an inertia constant, H, of four. The blue curve in Figure 10(a) shows the trajectory, calculated by the same simulation model but with all of the rotating machine inertia constants changed from 4.0 to 3.0 so that the rotating stored energy of the system is changed from 43200 MW-seconds to 32400 MW-seconds. The initial rate of decline of frequency is increased in accordance with the reduced inertia but the difference in the nadirs is much less than proportionate with the change in inertia.

Figure 10 (b) shows the initial decline of frequency after the loss of two percent of the total generation. The *initial* rate of decline of frequency in the two simulations corresponds closely to the values assigned to inertia, but this correspondence holds true only until change of turbine power becomes significant, and *is not* a useful indication of the level at which the decline of frequency will be arrested. Figure 10(c) now becomes useful; it shows the response over 20 seconds of one of the turbines making up the responsive part of the fleet, a conventional steam turbine in this instance. This response is quick, but certainly not instantaneous and is still in progress as the frequency reaches its nadir. The decline of frequency is arrested when the collective increase in turbine power is equal to the loss of generation. The timing and frequency at which the arrest is achieved are, therefore, as much dependent on the dynamics of turbine power variations as they are on system inertia.

The ability to maintain control of frequency after the initial arrest is as important as the arrest itself. The ability of the turbine(s) to increase their power quickly to arrest the frequency decline must be accompanied by the ability and willingness to sustain the increase. Figure 10(d) now becomes important. The red trace of Figure 10(d) shows the case where the primary response of the *responsive-non-sustaining* turbines is removed slowly; the recovery of frequency after the initial arrest is slow but positive. The black trace in Figure 10(d) shows the consequence of allowing secondary control to act too aggressively; the system is saved only by load shedding at 59.1Hz. (It is significant that the difference between the red and black traces is the withdrawal of the primary response of only ten percent of the responsive generation.)

9. Main points regarding control of frequency

9.1 Inertia

The base case example shows that, while the inertia of the system is the dominant influence on the *initial* rate of decline of frequency, it is far from a dominant influence on the frequency at which the decline is arrested. The point is emphasized by Figure 11 and Figure 12.

Figure 11 shows simulations of the system with the electronic generation fraction set to zero so that all its generation is rotating. The red trace shows the simulated response to the loss of 0.02 per unit generation when all rotating machines have their inertia constants set to 4.0 (red). The blue trace is the simulation result for the case with all machine inertia constants set to 3.0. Reducing the inertia constant from 4 to 3, thereby reducing the system stored energy from 48000MW-sec to 36000MW-sec, has a visible but not significant effect on the trajectory of frequency. Figure 12 shows simulations of the system with the electronic generation fraction set to 0.5. Replacing half of the rotating generation with electronically coupled generation reduces the system total inertia values to 24000MW-sec for the red traces and to 18000MW-sec for the blue traces. The frequency nadirs shown in Figure 11 and Figure 12 are summarized in Table 9. It shows that reducing the effective inertia constant of the system by a factor of 2.667 changes the magnitude of the frequency dip by a factor of only 1.27. The total inertia of the system is clearly not a good indicator of the magnitude of the frequency excursion that a given loss of generation will cause.

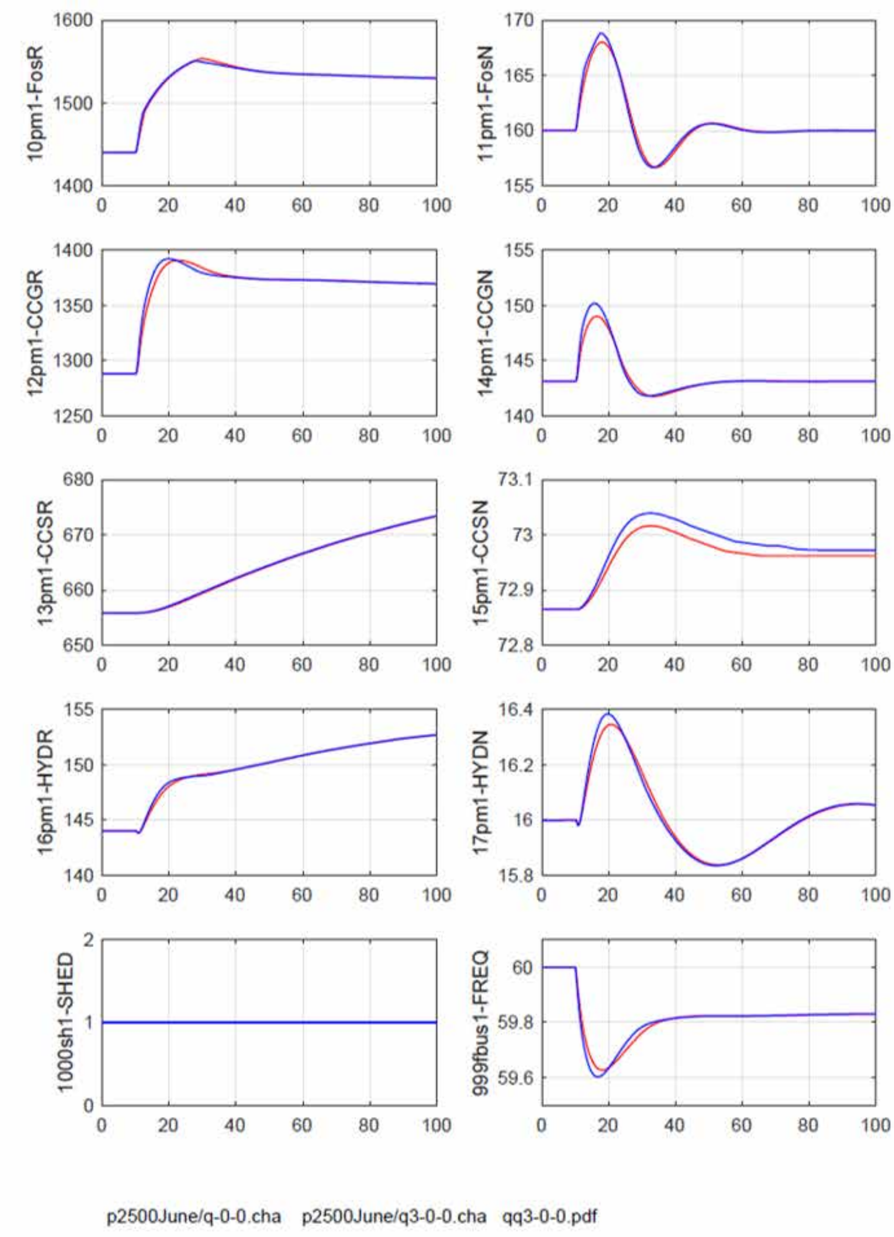

Figure 11. Effect of changing H constant of rotating fleet - Electronic fraction = 0.0

Red - Run q-0-0 - Efrac=0 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 - H=4
Blue - Run q3-0-0 - Efrac=0 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 - H=3

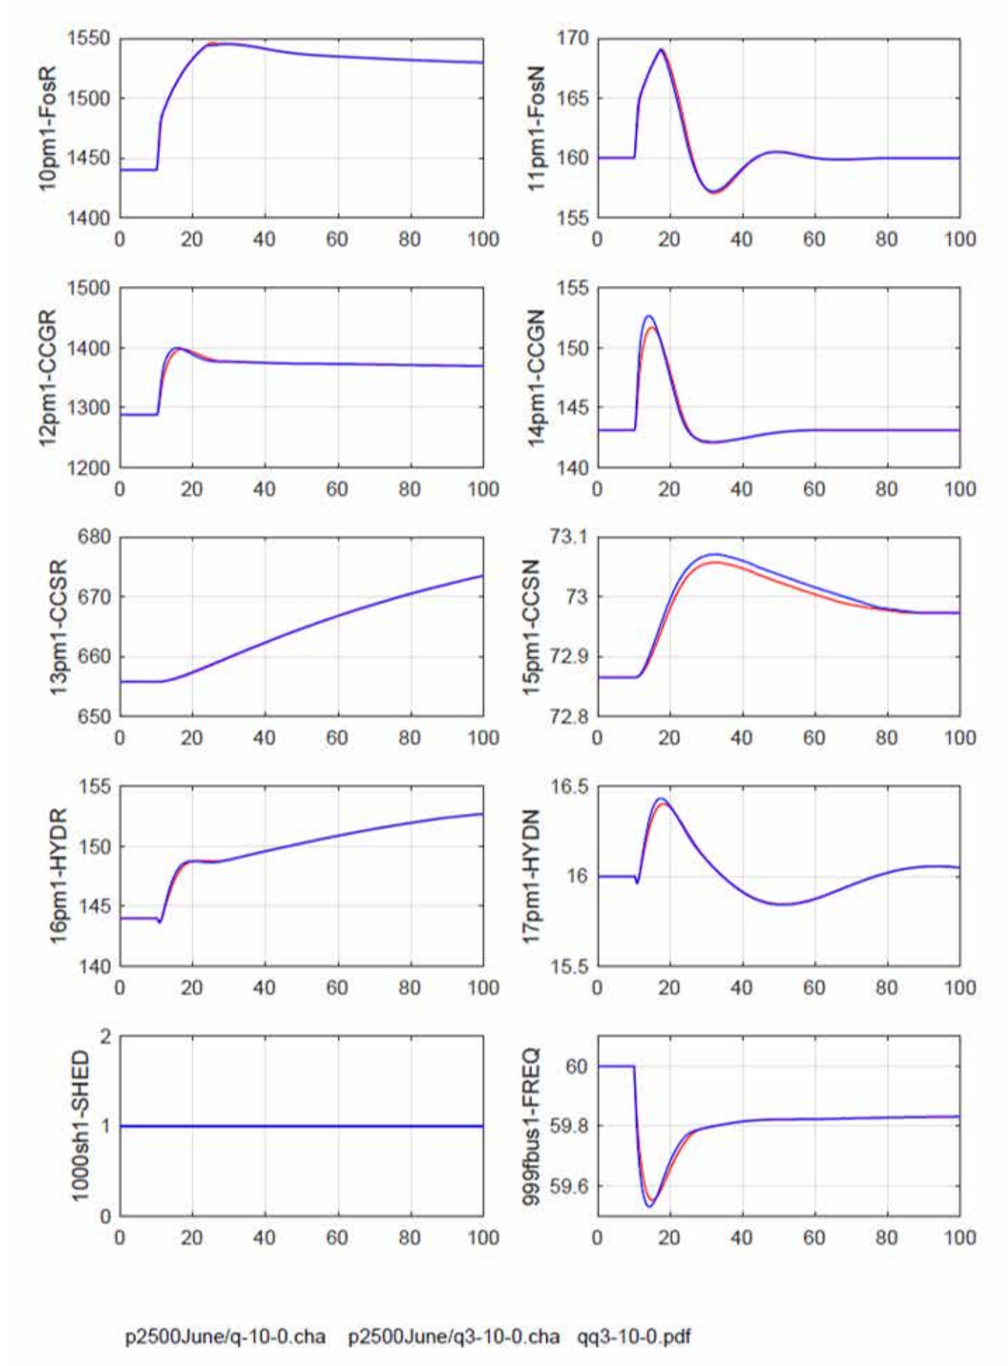

Figure 12. Effect of changing H constant of rotating fleet - Electronic fraction = 0.5

Red - Run q-10-0 - Efrac=0.5 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 - H=4
Blue - Run q3-10-0 - Efrac=0.5 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 - H=3

Table 9. Dependence of frequency nadir on system total inertia

Figure	Curve	Inertia (MW/s)	Effective System H	Nadir (Hz)	Δf (Hz)	Inertia Factor	Δf Factor
Figure 12	Red	48000	4.0	59.63	-0.37	1.000	1.000
Figure 12	Blue	36000	3.0	59.60	-0.40	1.333	1.081
Figure 13	Red	24000	2.0	59.55	-0.45	2.000	1.216
Figure 13	Blue	18000	1.5	59.53	-0.47	2.667	1.270

It should be noted that changing a half of the generation from rotating to electronic in these two examples was accomplished with no change in the amount of generation that is sensitive to change in frequency. (Whether this would be achieved in reality when replacing a half of a rotating fleet may be open to question. The generation that would be replaced would likely be older plants, but the age of the retired plants may not be a good indicator of whether they have been operated in frequency responsive mode.)

Figure 13 compares the traces from Figure 11 and Figure 12 for the case of H=4. The system stored energy is reduced from 480000MW-sec to 240000MW-sec by replacing rotating generation with electronic generation. Halving the total inertia of the system advances the time at which the frequency nadir occurs and lowers it, but certainly does not double the depth of the frequency dip.

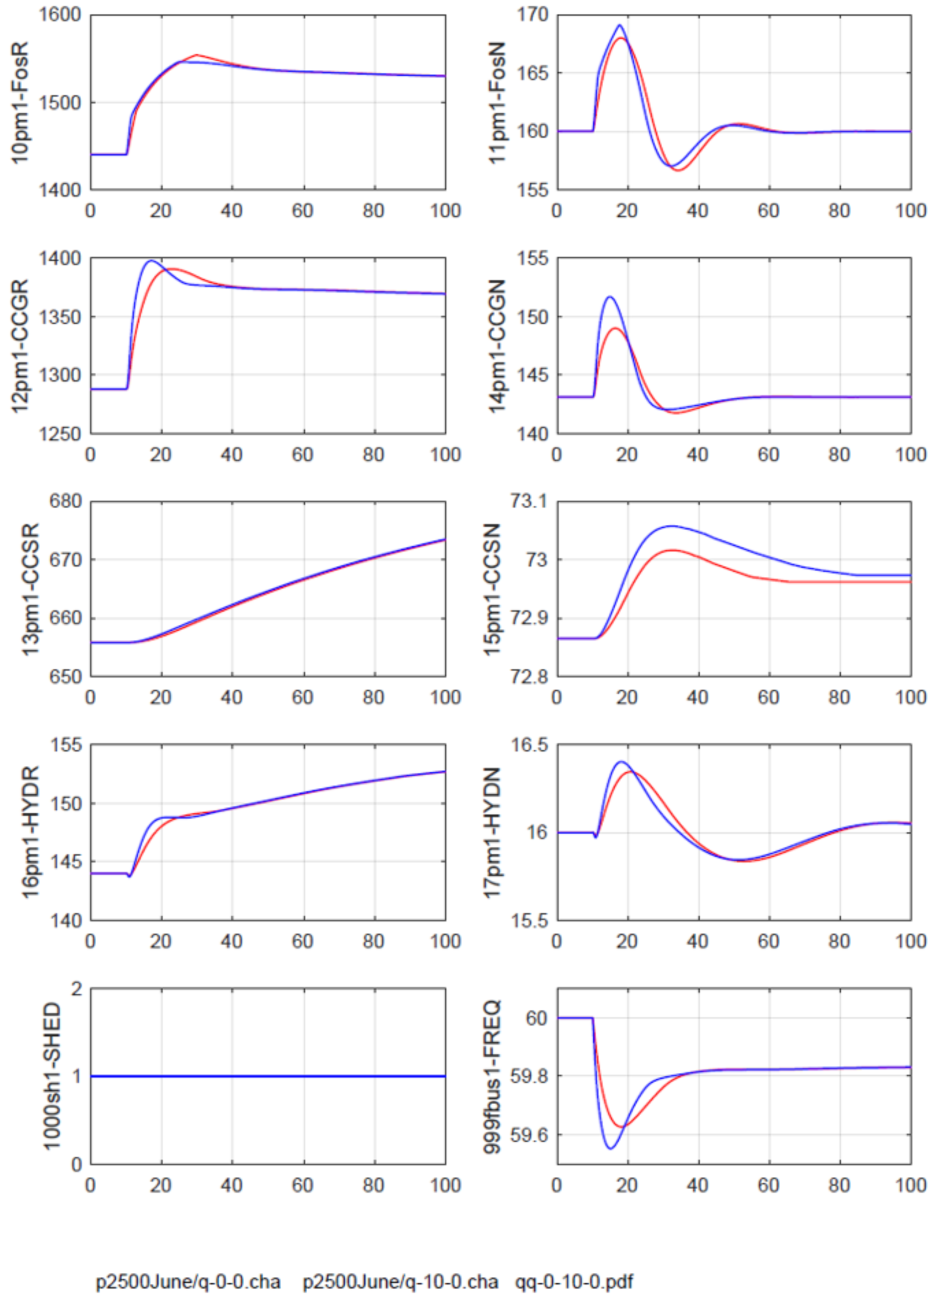

Figure 13. Effect of replacing non-responsive rotating generation with non-responsive electronic generation

Red - Run q-0-0 - Efrac=0.0 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 - H=4
Blue - Run q-10-0 - Efrac=0.5 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 - H=4

9.2 Speed of Governing Response

The speed of governing response is a strong influence on the nadir at which frequency is arrested. Figure 14 shows the frequency dips caused by the loss of 2 percent generation with well-tuned governing controls (red) and with the steam and gas turbine governors slightly detuned[4] (though still reasonably stable) to reduce their speed of response (blue). Comparing Figure 14 with Figure 11 makes it clear that the timing of governor response is as strong an influence on frequency nadir as is the inertia constant or total stored energy of the system.

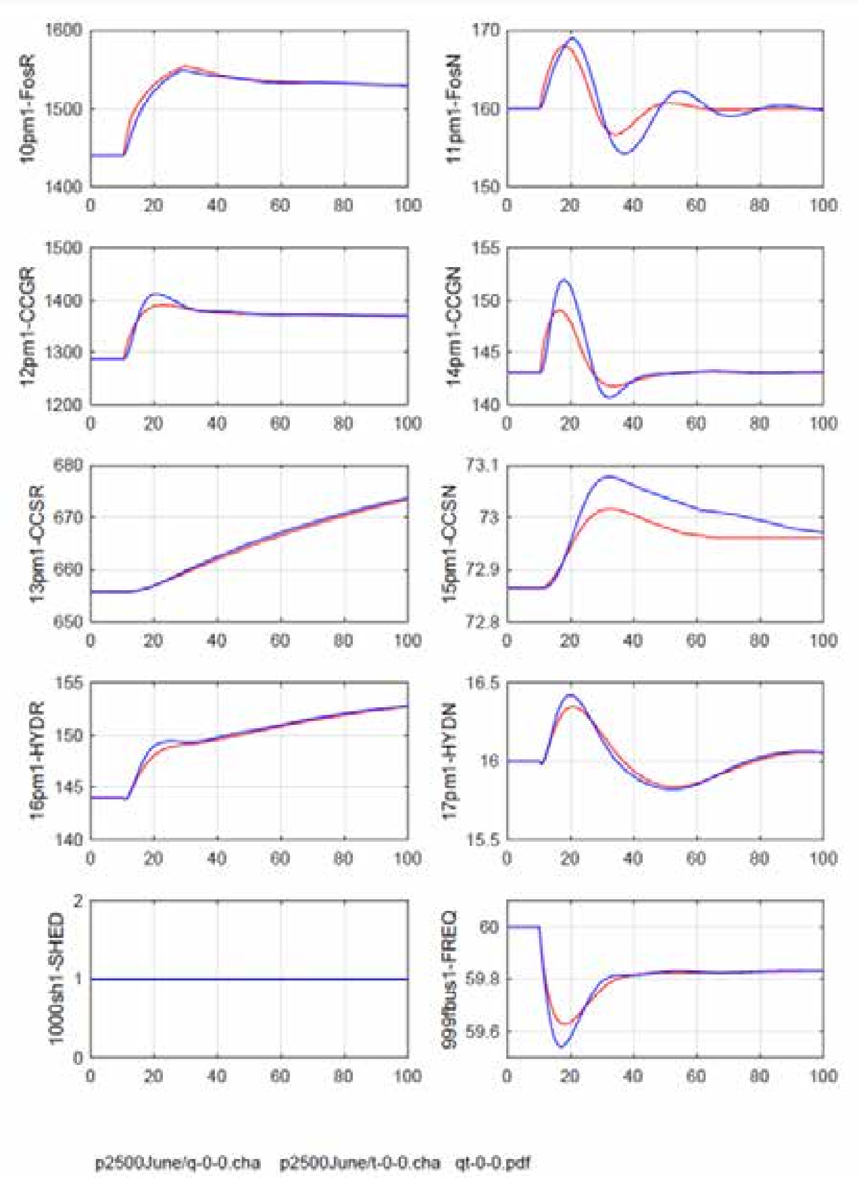

Figure 14. Standard and slow governing response

Red - Run q-0-0 - Efrac=0 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 – Normal thermal unit governing
Blue - Run t-0-0 - Efrac=0 - Rfrac=0.4 - Sfrac=0.9 - tact=2.0 – Detuned thermal unit governing

[4] By changing the valve actuator time constant, Tact, of the ggov1 controller model

9.3 Timing of Secondary Control Action

The speed of governing response is a strong influence on the frequency at which the decline of frequency is arrested but has a relatively minor influence on the recovery after the arrest. In contrast, the timing of secondary control action has only a minor effect on the initial arrest but a critical effect on the frequency trajectory after the initial arrest.

Four figures are needed to support this point. Figure 15 through Figure 18 show the effect of changing the rate of secondary control action by setting the controller gain, Ki, of the lcfb1 model to 0.002, 0.01, 0.015, and 0.02, respectively.

Figure 15 shows the behavior of an all-rotating system when the timing of the secondary control action is 'cautious' (red) and 'aggressive' (blue). The timing of the secondary control action is readily observable in the power trajectories in right-hand column of the figure, but has no significant effect on the frequency nadir. Figure 16 shows the corresponding behavior when 50 percent of the generation that is rotating but non-responsive is replaced by electronic generation. This replacement causes the arrest of frequency decline to occur at 59.55Hz, rather than 59.60Hz, but otherwise makes no significant change in the form of the system's response.

Figure 17 compares simulations made with the same secondary control timing as Figure 15 but with a lesser fraction of the responsive generation sustaining its initial response; (sf is reduced from 0.9 to 0.6). With 'cautious' secondary control frequency recovers, but with aggressive secondary control the arrest of frequency decline is not maintained.

Figure 18 corresponds to Figure 17 for the case with one half of the rotating fleet replaced by electronic generation. As above, the replacement of rotating generation did not change the amounts of responsive-sustaining and responsive-non-sustaining generation.

For the simulation shown by Figure 17 and Figure 18 the fraction of the generation that is frequency responsive and sustaining is 0.4*0.6=0.24. The increase in output required from each turbine in this fraction is 0.02/0.24=0.0833 per unit. The upward maneuvering range in this generation is 0.09 per unit and so the required amount of sustainable primary response would seem to be available. It is not all available with the same timing, however. The conventional steam turbines and gas turbines deliver their primary response promptly, but the response of the steam turbines in the combined cycles, while sure to be delivered, is delivered slowly. Withdrawing non-sustaining generation too quickly in relation to the rate at which primary response is delivered, leaves the system without sufficient generation; collapse can be avoided only by shedding load.[5]

[5] The behavior illustrated by Figure 17 and Figure 18 is not academic; it was observed at large-grid scale in Malaysia in 1996 and in Italy in 2003. It has also been observed in all-gas-turbine systems of medium size in the oil production industry.

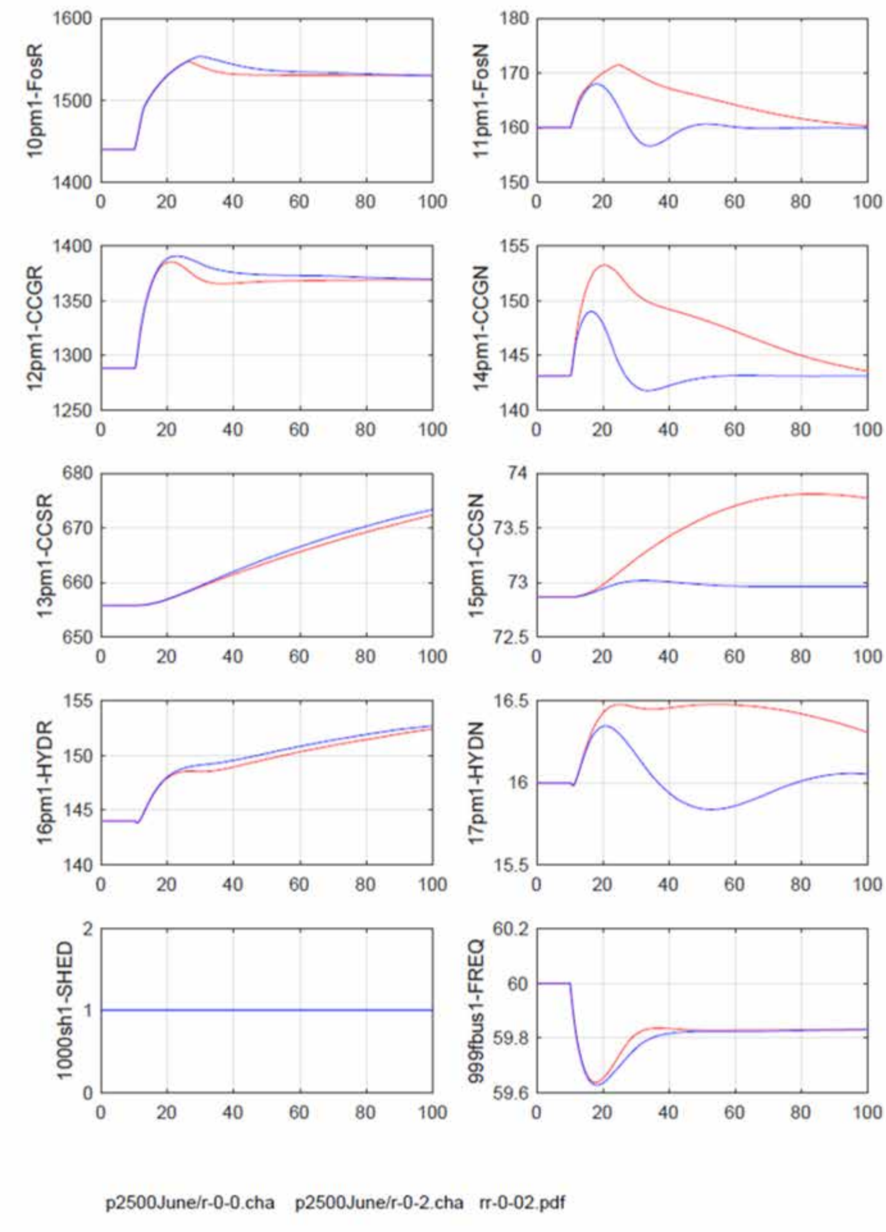

Figure 15. Cautious and Aggressive secondary control action – no electronic generation

Red - Run r-0-0 - Efrac=0 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 – Normal secondary control gain
Blue - Run r-0-2 - Efrac=0 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 – Aggressive secondary control gain

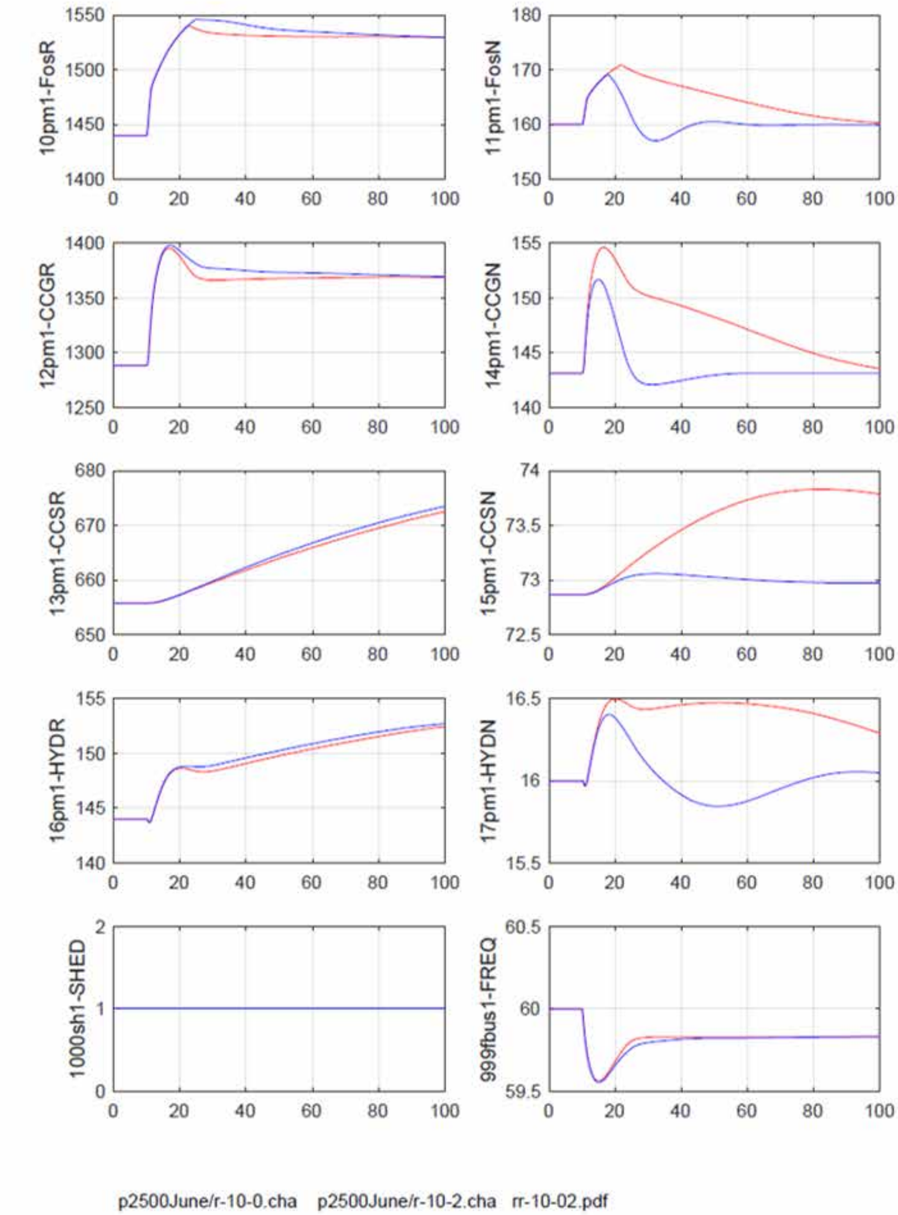

Figure 16. Cautious and Aggressive secondary control action – 50 percent electronic generation

Red - Run r-10-0 - Efrac=0.5 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 – Normal secondary control gain
Blue - Run r-10-2 - Efrac=0.5 - Rfrac=0.4 - Sfrac=0.9 - tact=0.2 – Aggressive secondary control gain

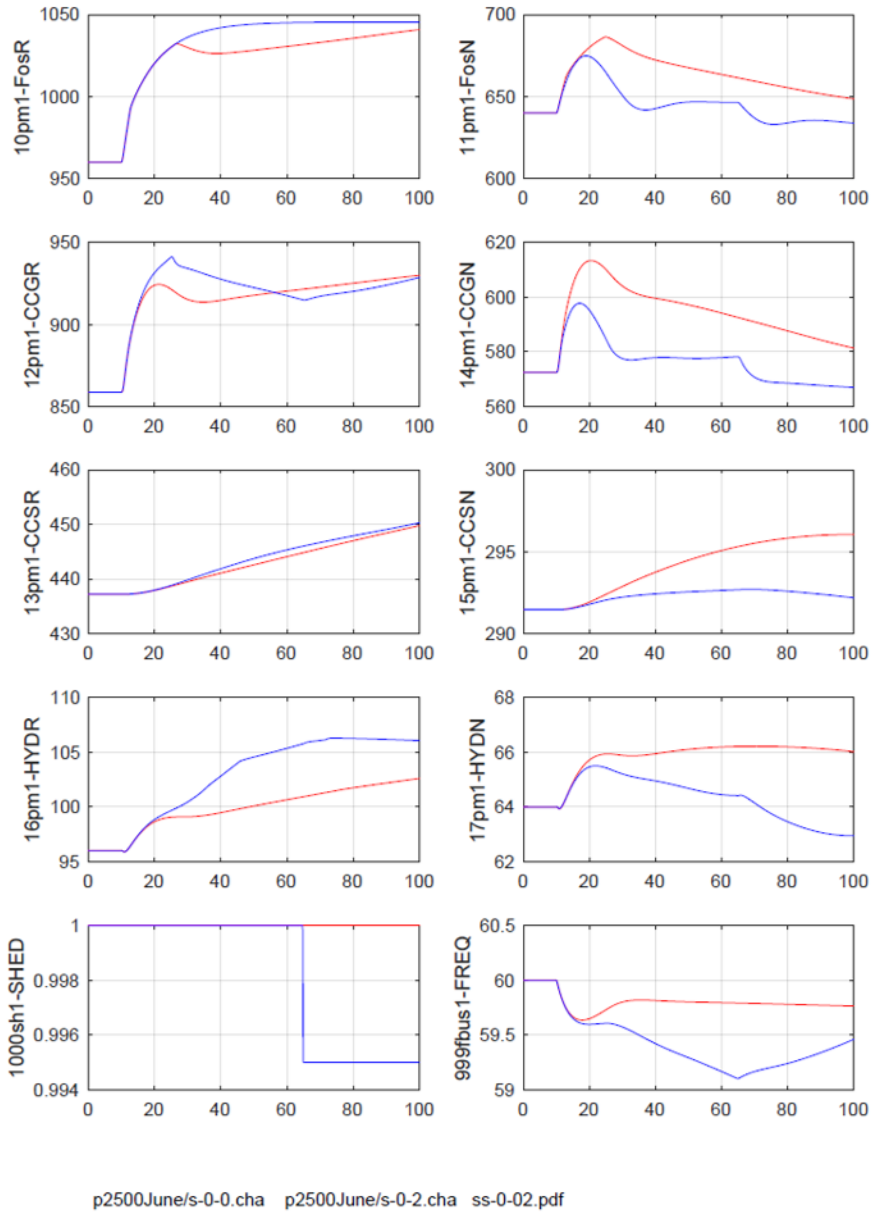

Figure 17. Cautious and Aggressive secondary control action – reduced sustaining fraction – no electronic generation

Red - Run s-0-0 - Efrac=0 - Rfrac=0.4 - Sfrac=0.6 - tact=0.2– Normal secondary control gain
Blue - Run s-0-2 - Efrac=0 - Rfrac=0.4 - Sfrac=0.6 - tact=0.2 – Aggressive secondary control gain

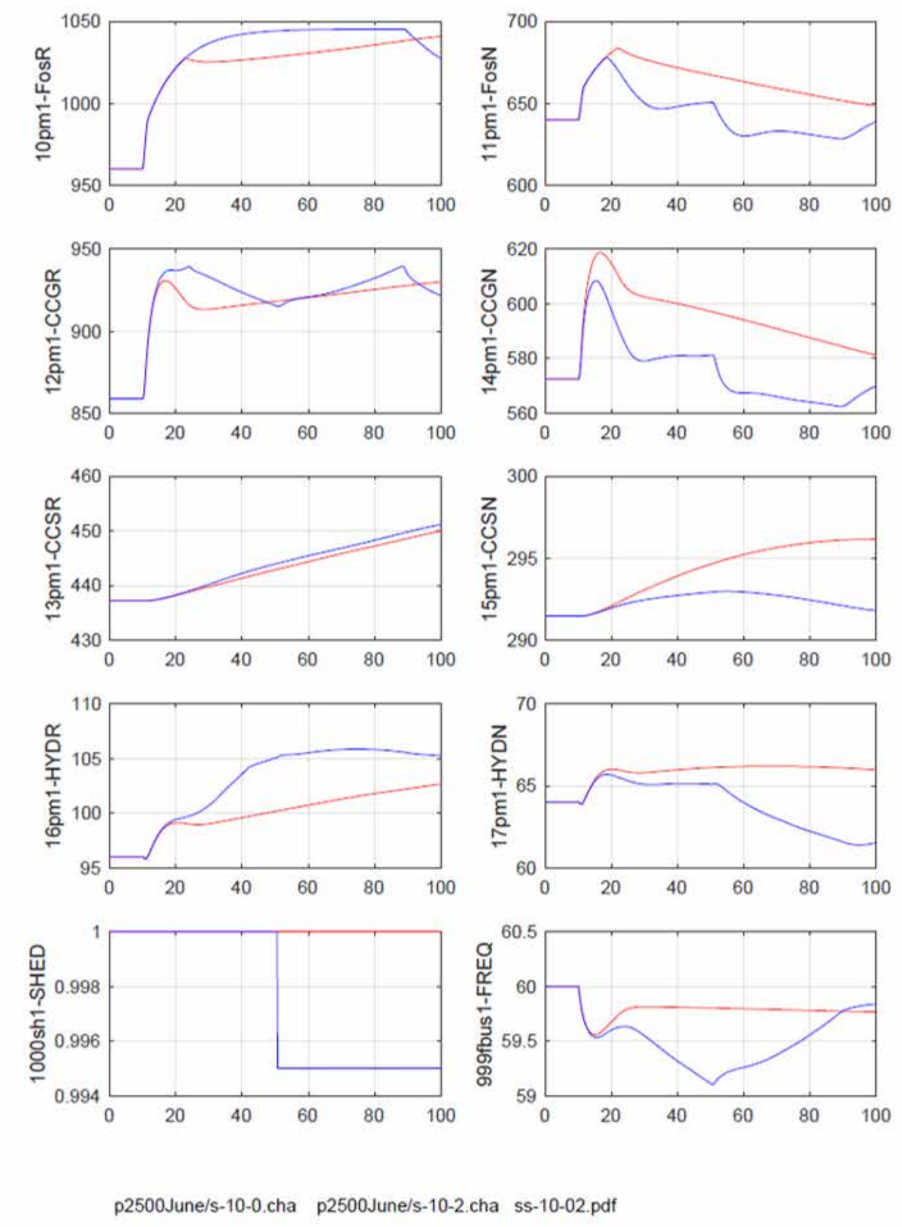

Figure 18. Cautious and Aggressive secondary control action – reduced sustaining fraction – 50 percent electronic generation

Red - Run s-10-0 - Efrac=0.5 - Rfrac=0.4 - Sfrac=0.6 - tact=0.2
Blue - Run s-10-2 - Efrac=0.5 - Rfrac=0.4 - Sfrac=0.6 - tact=0.2

9.4 System Behavior at Low Frequency

Figure 19 shows the behavior of system frequency and of a large gas turbine in a loss of generation simulation made with the same system modeling as used in the previous section, with the responsive-non-sustaining fraction chosen to cause the responsive-sustaining generation to go to maximum output.

The black curve shows the same initial arrest of frequency decline as seen in the base case example and the same partial recovery after the initial arrest. In this case, however the recovery fails and by 25 seconds after the initial loss, frequency is on an unrecoverable slide downwards.

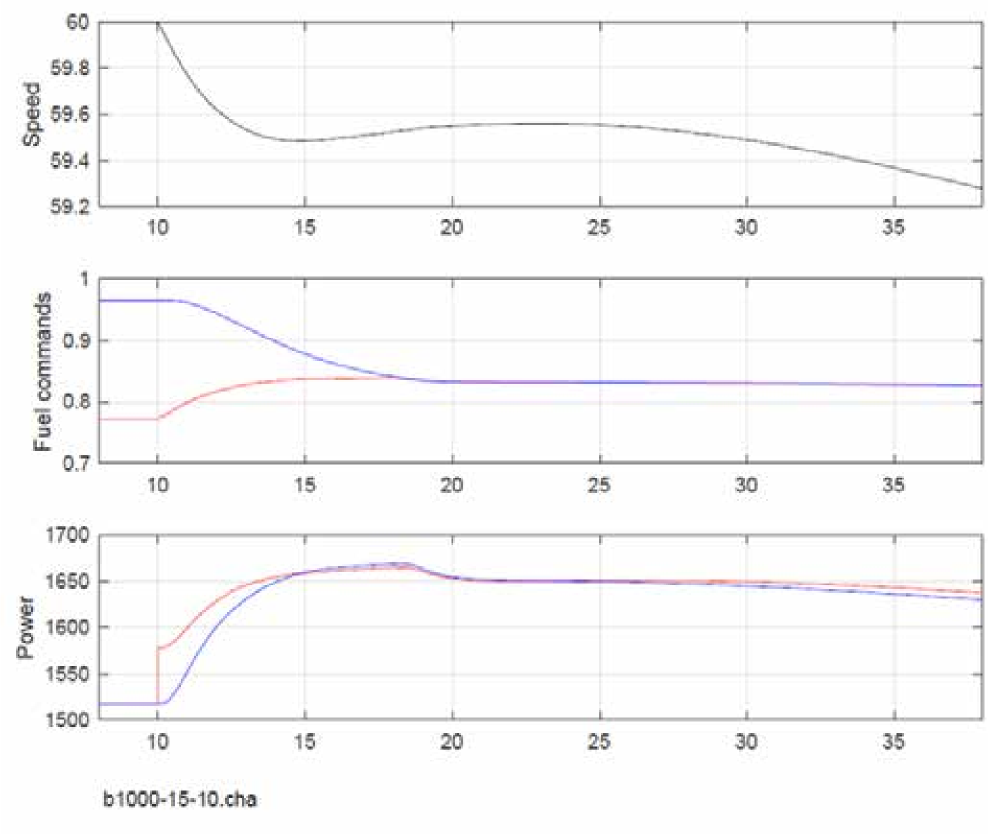

Figure 19. Response of a large single-shaft gas turbine in a severe loss-of-generation event

Middle – red – governor fuel command
Middle – blue – temperature limit fuel command
Bottom – red – generator electrical power
Bottom – blue – turbine shaft power

The behavior shown by Figure 19 is the result of these factors:

- Because only a fraction of the total generation is responsive to change of frequency, the increase in output required from the responsive generation is a substantially greater fraction of that generating capacity than the initial two percent loss is of the total generation.
- While all of the responsive-and-sustaining turbines have a margin of 9 percent of their capacity,

the sum of these margins on this subset of turbines is only very slightly greater than the two percent needed to make up the generation loss.

- The maximum outputs of the single shaft gas turbines decrease roughly as the square of their speed; at 59.5Hz these engines can deliver only approximately 98.3 percent of what they can deliver at rated speed. The total sustainable control response is reduced accordingly.

The combined effects of asking a few machines to produce a large response, of the withdrawal of a significant part of the initial primary response, and of the reduction in maximum output, is that by 25 seconds into the event there is not enough generating capability to meet the total load. The insufficiency at 25 seconds is small but, rather than decreasing asymptotically with time, it increases as declining speed reduces the maximum outputs of large gas turbines.

9.5 Recognition of Frequency Nadir

For this report simulations have been carried out until it is clear that either, frequency is stabilized and can be maintained, or frequency has fallen below 58.8Hz. The nadir of frequency is taken to be the lowest value seen; it is either a low value from which there is secure recovery or it is the low value, 58.8Hz, at which the simulation was terminated. A temporary nadir that is followed by a recovery and subsequent decline to a lower value is not recognized.

10. Parametric Examination of Primary Control Requirement

10.1 Key Parameters

The key parameters describing the amount of reserve for primary control are:

Rfrac: the fraction of the total generation that is responsive to change of frequency
Sfrac: the fraction of the responsive generation that sustains its primary control response

The fraction of total generation that responds to change of frequency and sustains its response is *Rfrac∗Sfrac* which is plotted as a hyperbolic surface in Figure 20.

If lost generation could be replaced instantaneously, the primary regulating reserve levels indicated by Figure 20 would be sufficient to enable the system to retain control of frequency after a sudden loss. Instantaneous response capability is a 'nice idea, but not realistic'. Battery based generation that can increase its output very nearly instantaneously is becoming available and may be a significant contributor to primary response in the foreseeable future. Given practical rates of primary response of rotating generation, a sustainable reserve margin indicated by Figure 20 to be equal to the anticipated loss of generation is not sufficient.

To estimate the level of sustainable regulating margin that must be maintained it is necessary to track the behavior of the system in the manner indicated by Figure 15 for the broad range of the key parameters, *Rfrac* and *Sfrac*.

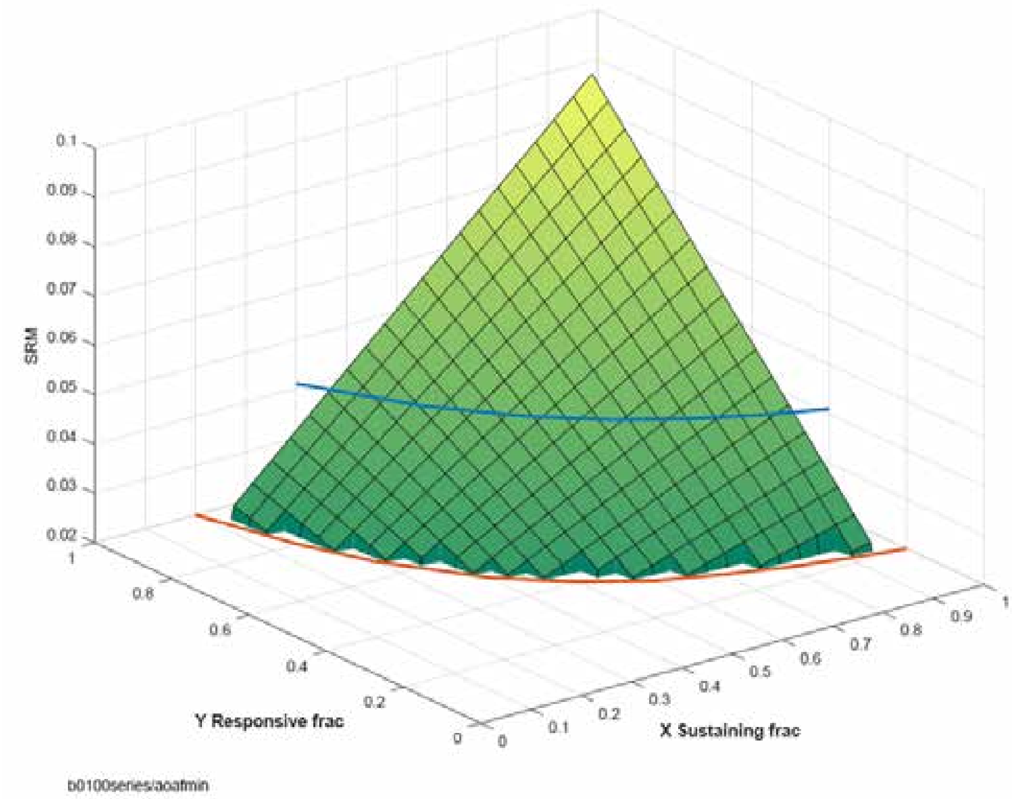

Figure 20. Available sustainable response capability

10.2 Mapping response

A useful characterization of the response of the system on the basis of *Rfrac* and *Sfrac* is a map of the nadir of frequency. Representative such maps are shown in Figure 21.

Figure 21 shows how the nadir of frequency depends on the fractions of the total generation that are responsive and sustaining. The Y axis of the plot gives the fraction that is responsive, *Rfrac*. The X axis gives the fraction, *Sfrac*, of the responsive generation that sustains its initial changes. The loss of generation is 4.0 percent of the total. The relatively flat top of the surface plot shows the nadir of frequency in cases where the reserve margin is sufficient so that:

- The responding-and-sustaining generation does not need to go to its maximum.
- The decline of frequency is arrested at a level such that the initial recovery is sustainable in the long term.
- Frequency does not go below 59.1 Hz.

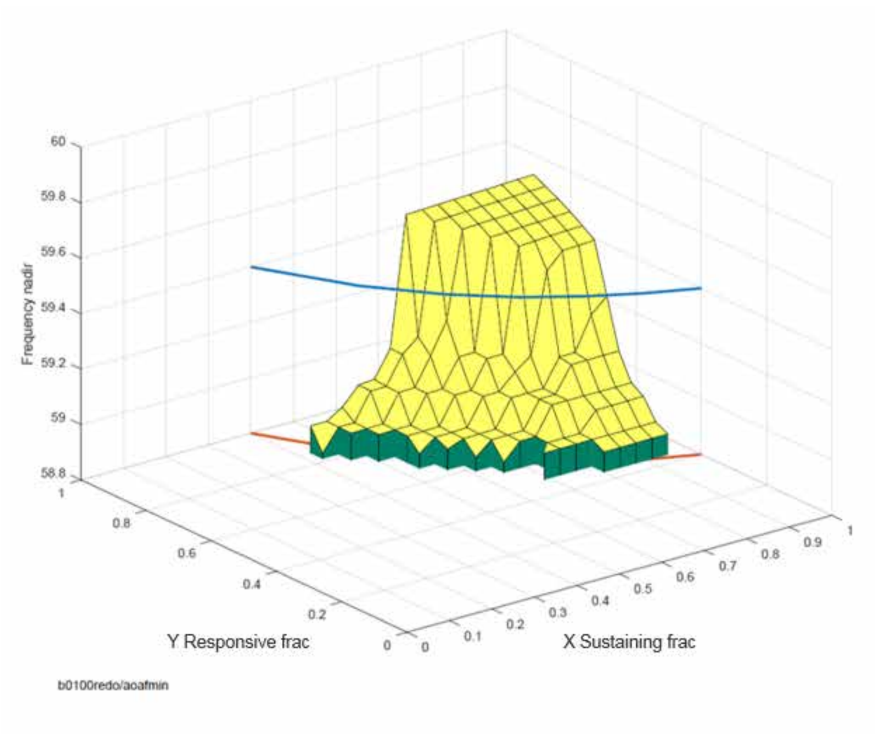

(a) 1 percent electronic

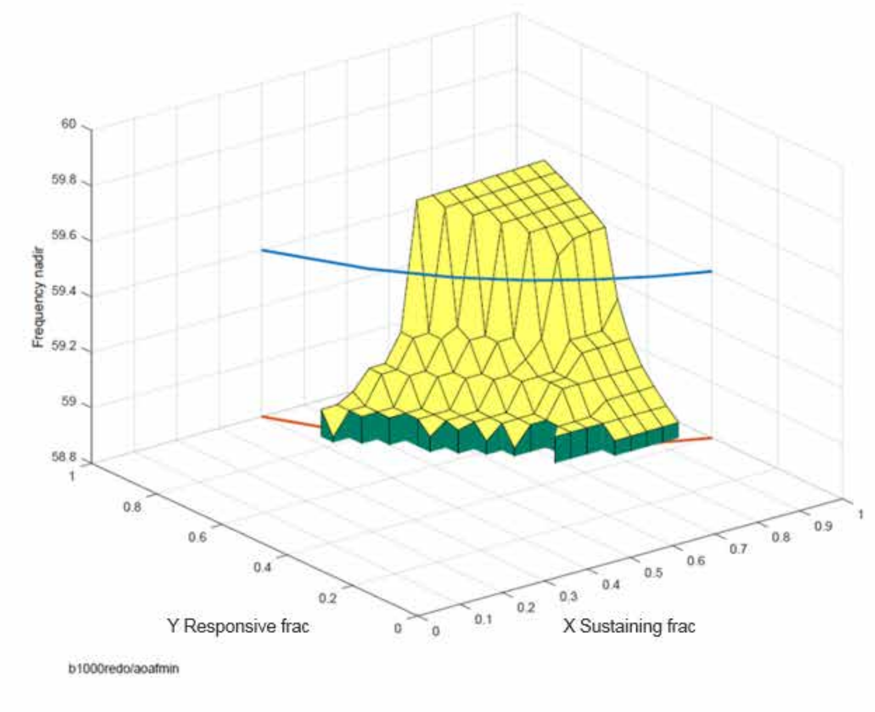

(b) 10 percent electronic

Figure 21. Frequency nadir map: thermal fleet

Red, Blue - Sustainable reserve margin = 0.04 per unit

The steep edge of the surface plot indicates that either the initial decline was not arrested above 58.8Hz or that recovery after an initial arrest was not maintained and frequency resumed a decline to 58.8Hz. The blue and red hyperbolas show the boundary at which $Rfrac * Sfrac = 0.04$ and the amount of responsive-sustaining generation is just equal to the loss of generation.

This plot effectively relates the ability to control frequency to the responsive and sustaining generation fractions in the following terms:

- With fractions corresponding to the upper part of the plot, the primary control capability of the system is sufficient with regard to both amount and dynamic characteristics.
- With fractions outside the steep edge of the plot, the primary control capability of the system is insufficient.

10.3 Four percent loss of generation

Figure 21 (top) describes the behavior of the system with minimal (0.01 per unit) electronic generation. Figure 21 (bottom) applies to the same system and same 4.0 percent generation loss but has the electronic generation as 0.1 per unit of the total.

Figure 21 shows that the fractions of responsive and sustaining generation would have to be in the high part of the range of present practice to be sure of riding through a four percent loss without load shedding. The difference between the maps in the two figures is visible (upon close examination) but not large, showing that using non-responsive electronically coupled generation to replace ten percent of the rotating fleet, from its non-responsive part, does not have a significant effect on the control of frequency. The loss of four percent of total generation is a very severe disturbance.

10.4 Two Percent Loss of Generation: Thermal Generating Fleet

Figure 22 through Figure 24 show surface plots of the frequency nadirs produced by a 2.0 percent generation loss for the case of a thermal generating fleet. The relative makeup of the rotating generation fleet is the same as considered in Section 10.3 and is shown in Table 2.

Figure 22 shows the map for the case with ten percent of the total in electronic generation. The cliff of the plot indicates that control of frequency could be maintained with as little as one half of the total generation being responsive, provided that at least a half of this responsive generation will sustain its initial governing response.

Increasing the fraction of electronic generation to 25 percent is shown by Figure 23 to require somewhat greater fractions of the generation to respond to frequency and sustain the response. The required increases would not call for major changes in operating practice, however.

Figure 24 shows the situation to be still more demanding when the fraction of electronic generation is increased to 40 percent of the total. With 40 percent of the total being electronic the maximum fraction of generation that can be responsive to frequency is 60 percent and, recognizing that there will surely be some generation that cannot respond, the highest practical level will be less than 60 percent.

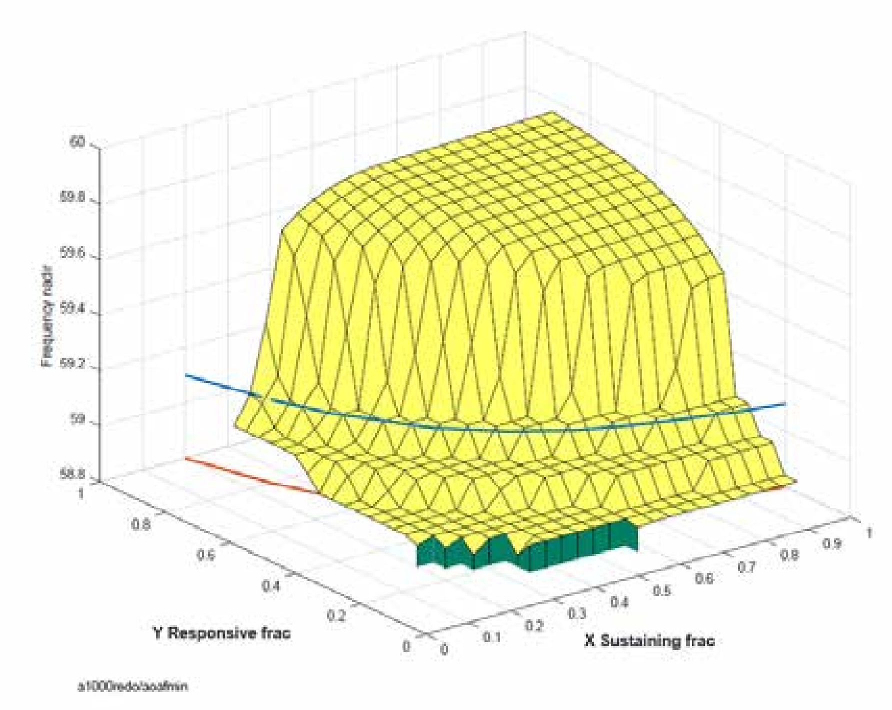

Figure 22. Frequency nadir map: 10% electronic generation (a1000), thermal fleet

Blue - Sustainable reserve margin = 0.02 per unit

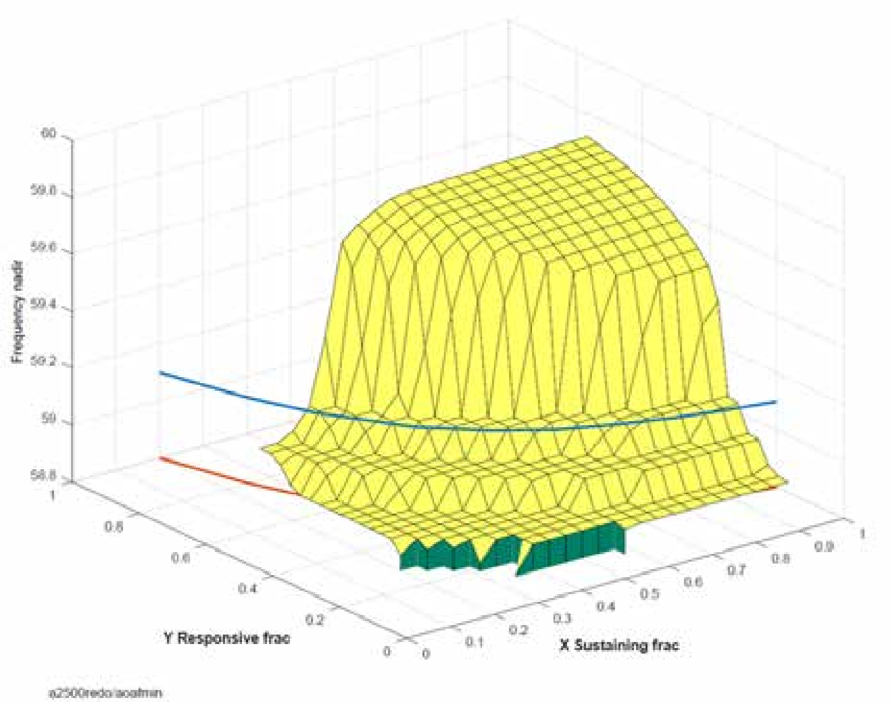

Figure 23. Frequency nadir map: 25% electronic generation (a2500), thermal fleet

Blue - Sustainable reserve margin = 0.02 per unit

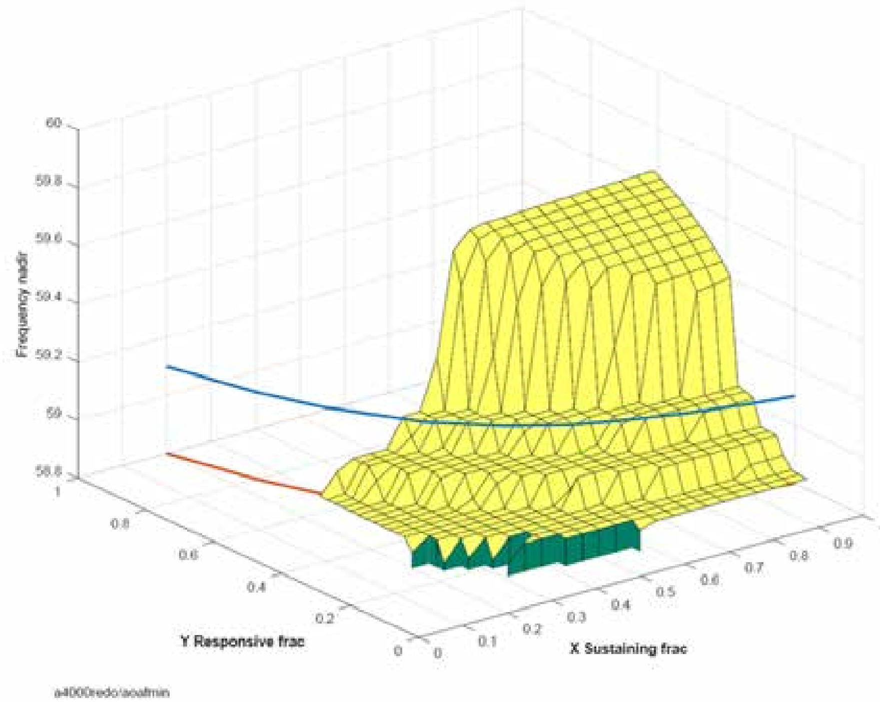

Figure 24. Frequency nadir map: 40% electronic generation (a4000), thermal fleet

Blue - Sustainable reserve margin = 0.02 per unit

10.5 Two Percent Loss of Generation: Mixed Generating Fleet

Figure 25 through Figure 27 are surface plots of the frequency nadirs shown by simulations of a 2 percent generation loss in a system with the mixed thermal-hydro generation fleet whose makeup is given in Table 3.

Figure 28 shows the response surfaces for case of 10 percent electronic generation with the thermal rotating fleet (top) and mixed thermal-hydro fleet (bottom). The 'cliff' indicating the minimum workable responsive generation fraction has moved (requiring a greater responsive fraction with the mixed fleet) but the difference is not extreme.

Figure 29 and Figure 30 make the comparison for the cases with 25 and 40 percent electronic generation. These two figures indicate that the requirements of the thermal and mixed-hydro fleets for responsive generating capacity differ in detail but are the same in principle.

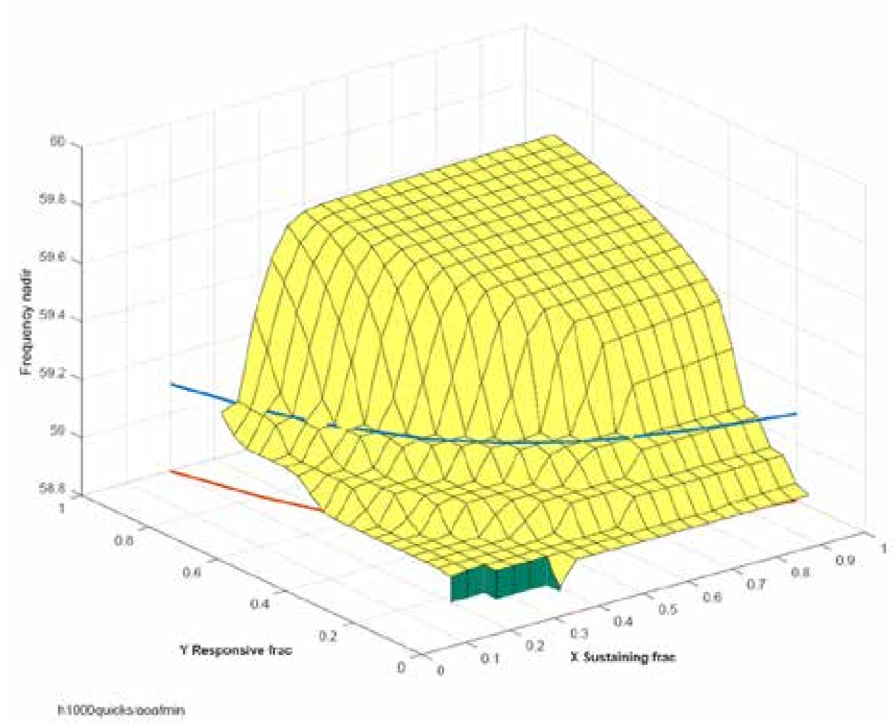

Figure 25. Frequency nadir map: 10% electronic generation (h1000quicks), mixed thermal-hydro fleet

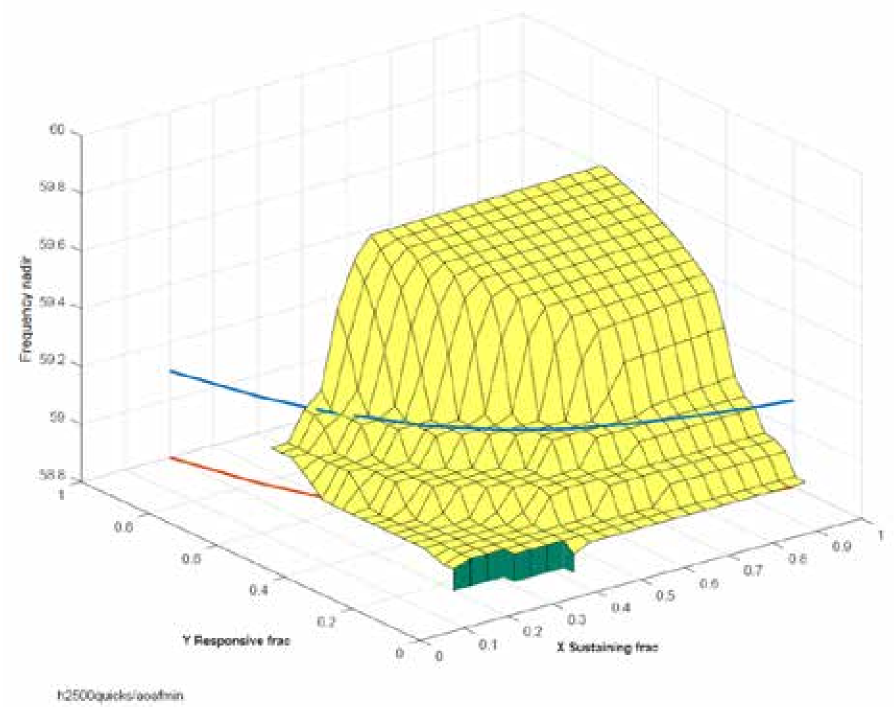

Figure 26. Frequency nadir map: 25% electronic generation (h2500quicks), mixed thermal-hydro fleet

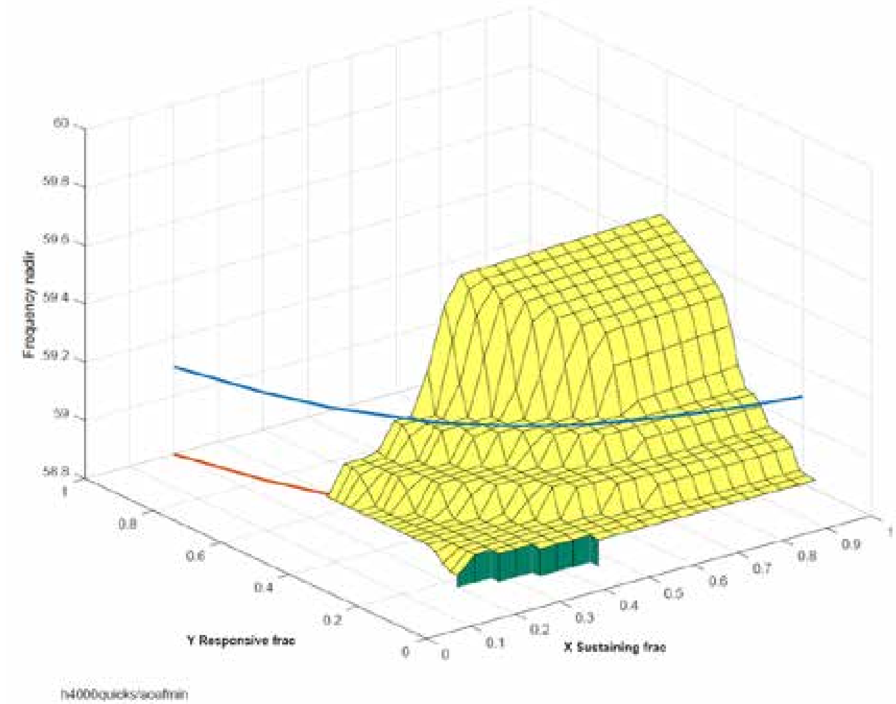

Figure 27. Frequency nadir map: 40% electronic generation (h4000quicks), mixed thermal-hydro fleet

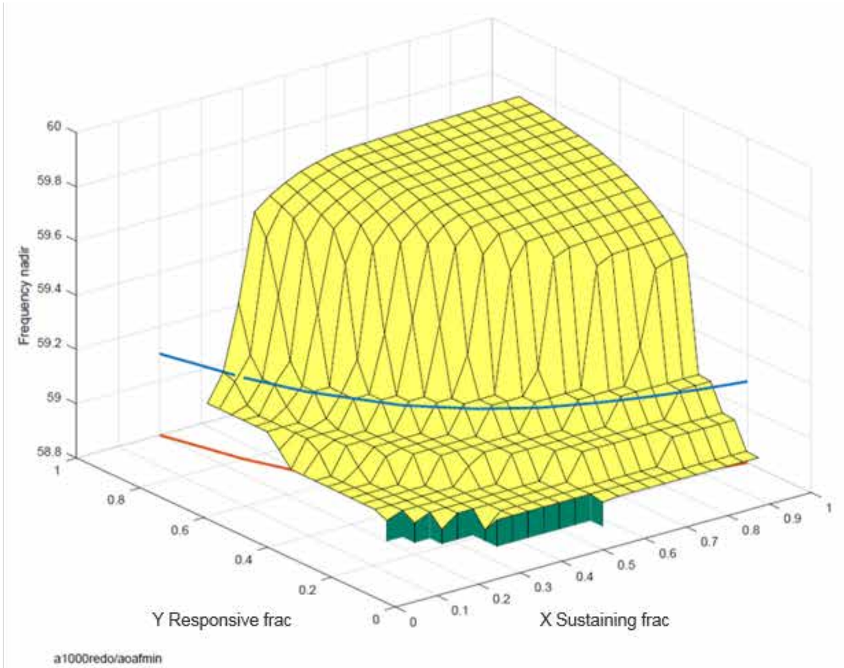

(a) Thermal fleet

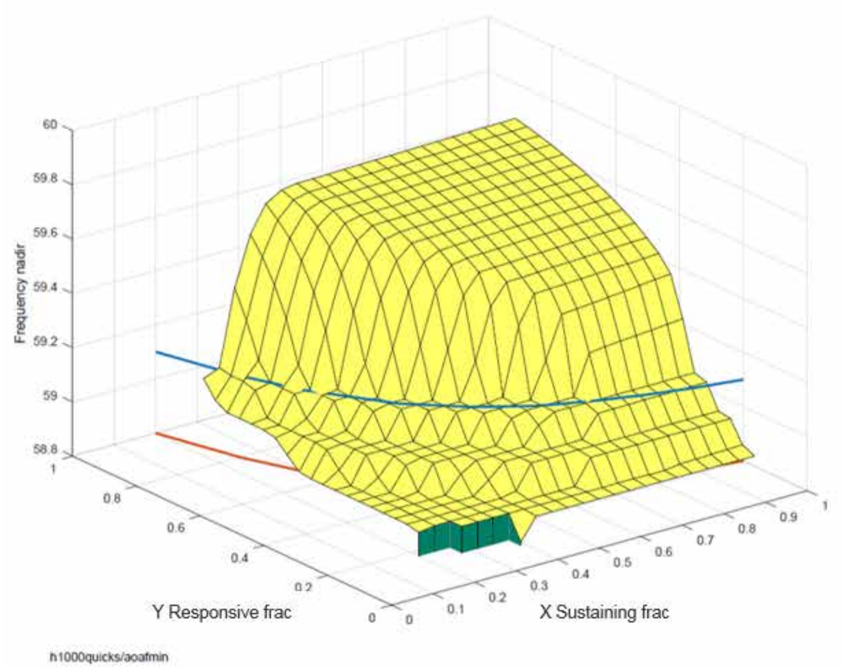

(b) Mixed thermal-hydro fleet

Figure 28. Comparison of response with thermal and mixed rotating fleet: 10% electronic generation

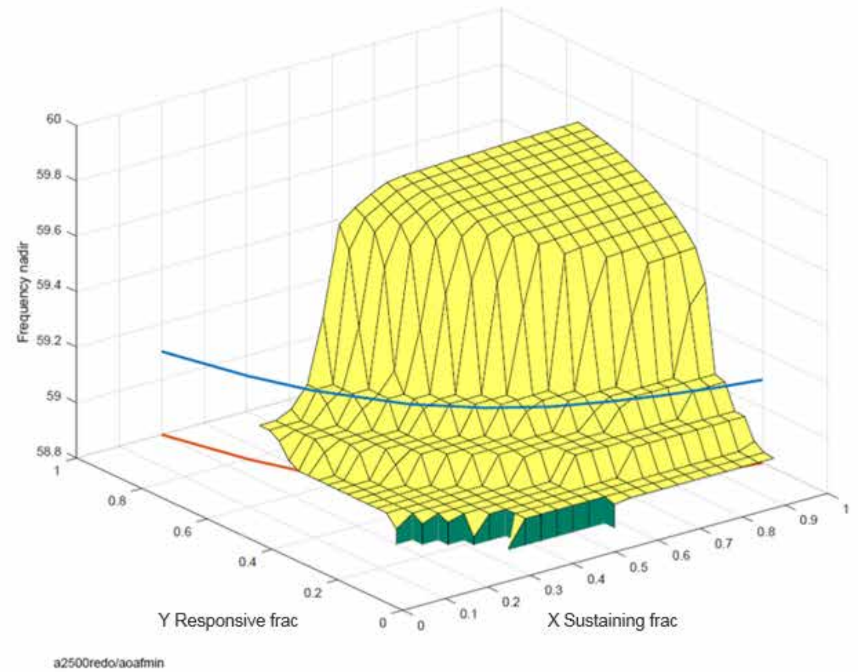

(a) Thermal fleet

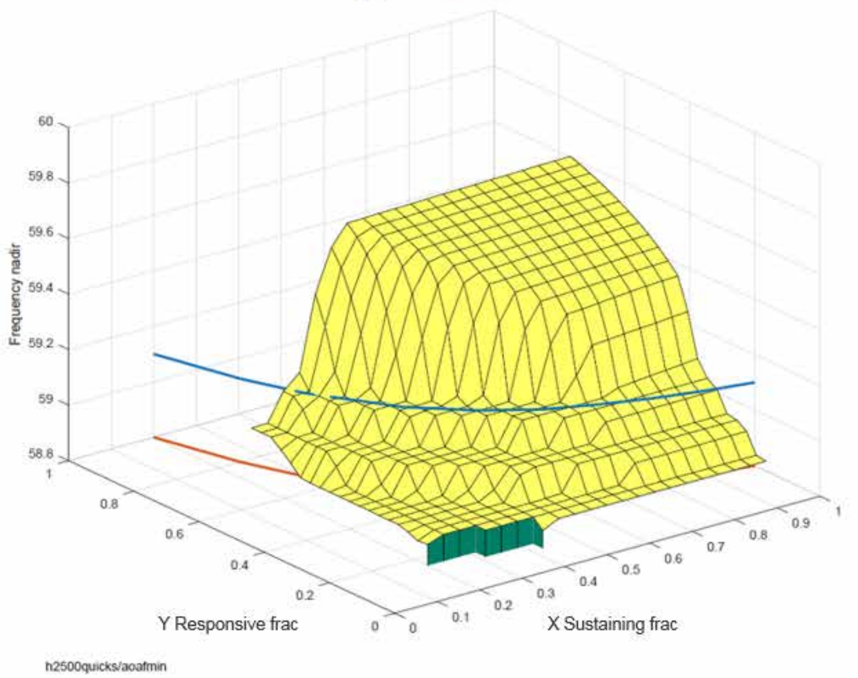

(b) Mixed thermal-hydro fleet

Figure 29. Comparison of response with thermal and mixed rotating fleet: 25% electronic generation

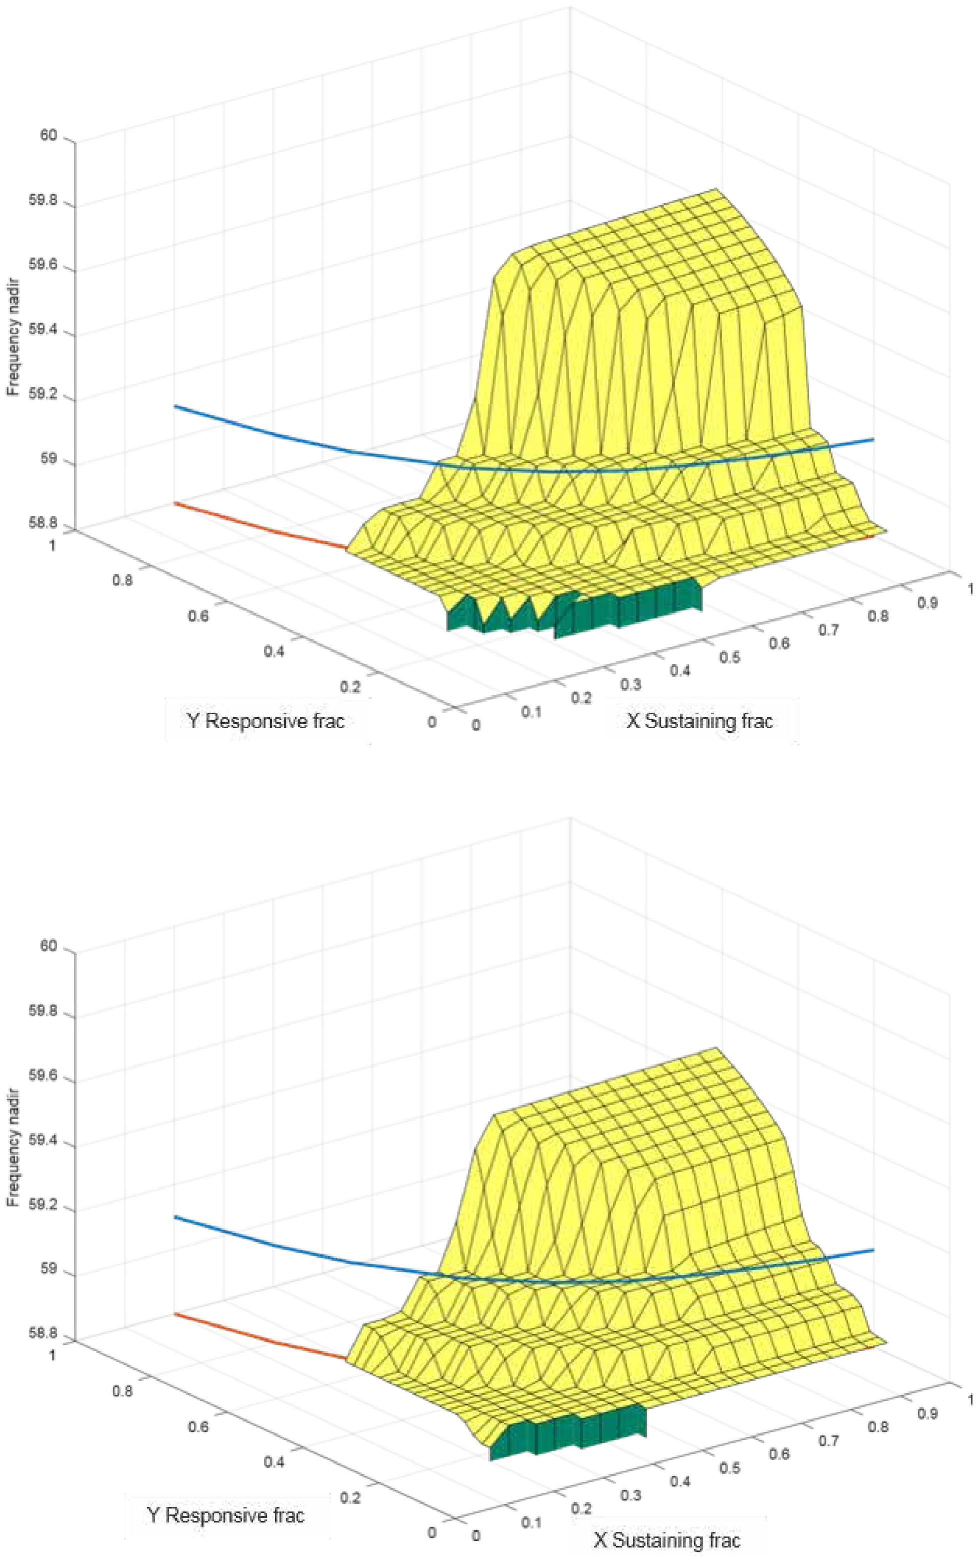

Figure 30. Comparison of response with thermal and mixed rotating fleet: 40% electronic generation

11. Power Plant Load Control with Frequency Bias

11.1 Power Plant Operating Modes

The withdrawal of turbine power increases by power plant load controllers after the arrest of a frequency dip is detrimental to the control of grid frequency. This detrimental effect can be alleviated by:

- Redefining the scheduled power target used in a power plant to mean the power that is to be achieved when frequency is at its scheduled value (nominally 60Hz, but adjusted from time to time to achieve time error correction); and
- Applying a frequency bias to the scheduled power so that the power target for the plant controls (governors, boiler controls, and others as applicable) becomes

$$P_{set}(f) = P_{sched} + B_{fp}\,(f_{sched} - f)$$

Equation 5

Note that the gain, B_{fp}, is a parameter of the power plant secondary control and is not the same thing as the frequency bias (often denoted by B or B_f) used in discussion of grid-level balancing controls.

When operating with a frequency bias in effect, in accordance with Equation 5, a power plant delivers its scheduled power when grid frequency is normal and, 'as a good citizen' sustains its contribution to the correction of a persistent deviation of frequency from nominal.

As noted in Section 5, most power plants are operated in pure droop mode, prescheduled output mode without frequency bias, or prescheduled output mode with frequency bias.

Pure droop control is found in older power plants that have not been brought up to date. It is also an option in the digital control systems (DCS) found in most modern power plants. While it is widely available it is not widely used in U.S. power plants, largely because there is no clear administrative requirement for it to be used.

Prescheduled output mode without frequency bias is the prevalent mode of operation in U.S. power plants. It is consistent with power trading and scheduling practices and is seen by control room operators as the automation of what they are employed to do. Further, because there is widespread belief that variations of output incur 'wear and tear' on the plant, it is often used with the argument that it is beneficial to the reliability of the plant.

Prescheduled output mode with frequency bias is not widely used in U.S. power plants but is used by 'command' of the transmission system operator in many nations whose grids are fully developed (e.g., The United Kingdom, Malaysia, Italy, and Ireland). (Power plants in the ERCOT system of Texas do use frequency bias in their plant-level controls.)

11.2 Implementation of Frequency Bias

The manner in which a frequency bias is applied in a plant depends on its type and on configuration of controls in addition to its turbine control. In a straightforward gas turbine plant, load control with frequency bias could be achieved by control logic similar to the representative configuration shown in Figure 31.

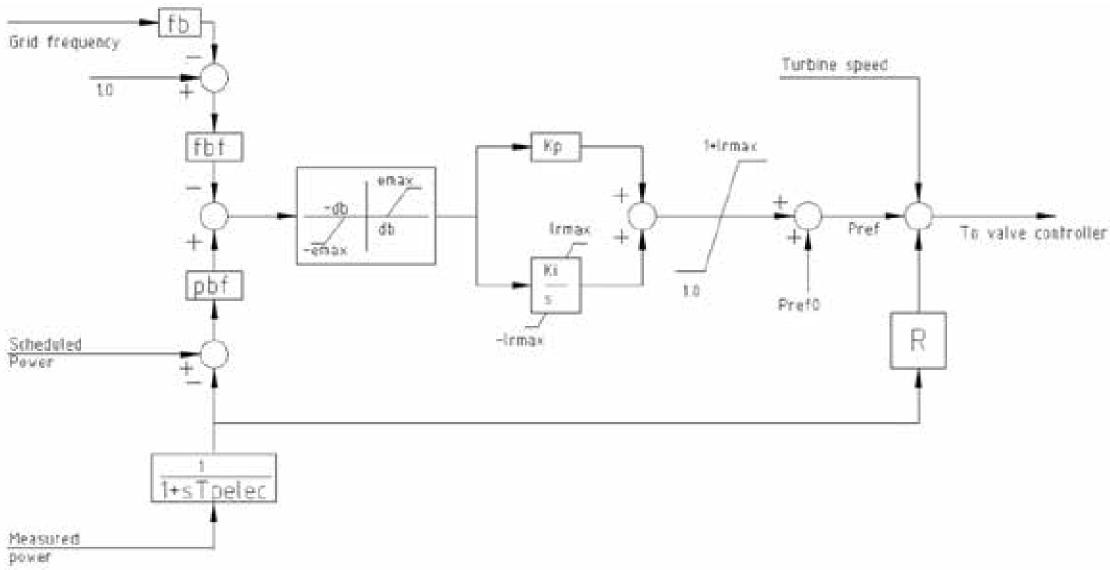

Figure 31. Representative load controller configuration

11.3 Illustrative Simulation

Figure 32 shows the effect of applying frequency bias to all of the responsive/non-sustaining turbines. Forty percent of the generation fleet is responsive, *Rfrac*=0.4 and fifty percent of the responsive generation is responsive-sustaining. The application of frequency bias transforms the responsive/non-sustaining generation, partly or fully depending on the bias factor, into responsive-sustaining generation.

The responsive-sustaining generation blocks (10, 12) are operating in pure droop mode with no adjustment of their speed-load references. The responsive/non-sustaining blocks (11, 14) have their speed-load references adjusted by the plant secondary controls. Because *Sfrac*=0.5 the size and initial outputs of blocks 10, 11 are equal and an ideal application of frequency bias should cause them to settle with equal increases in output. The same equality should apply to the two gas turbine blocks (12, 14).

The red traces show the response in the absence of frequency bias. Compare the power responses of the two steam turbine generation blocks, 10 and 11. Block 10 sustains its output increase while block 11 returns almost fully to its initial output. The same relative behavior is seen in the red traces for the large gas turbine blocks (12, 14).

The green traces show the response when the bias factor is set to $1/2R$ so that it is equivalent to a droop of twice the governing droop. Now the output of generation blocks 11 and 14 is still withdrawn after the initial response but to a lesser extent than without frequency bias.

The blue traces show the response when the bias factor is set to $1/R$ so that it corresponds exactly to the governing droop. With this bias the responses of generation blocks 10 and 11 settle with equal increases in output, as do blocks 12 and 14.

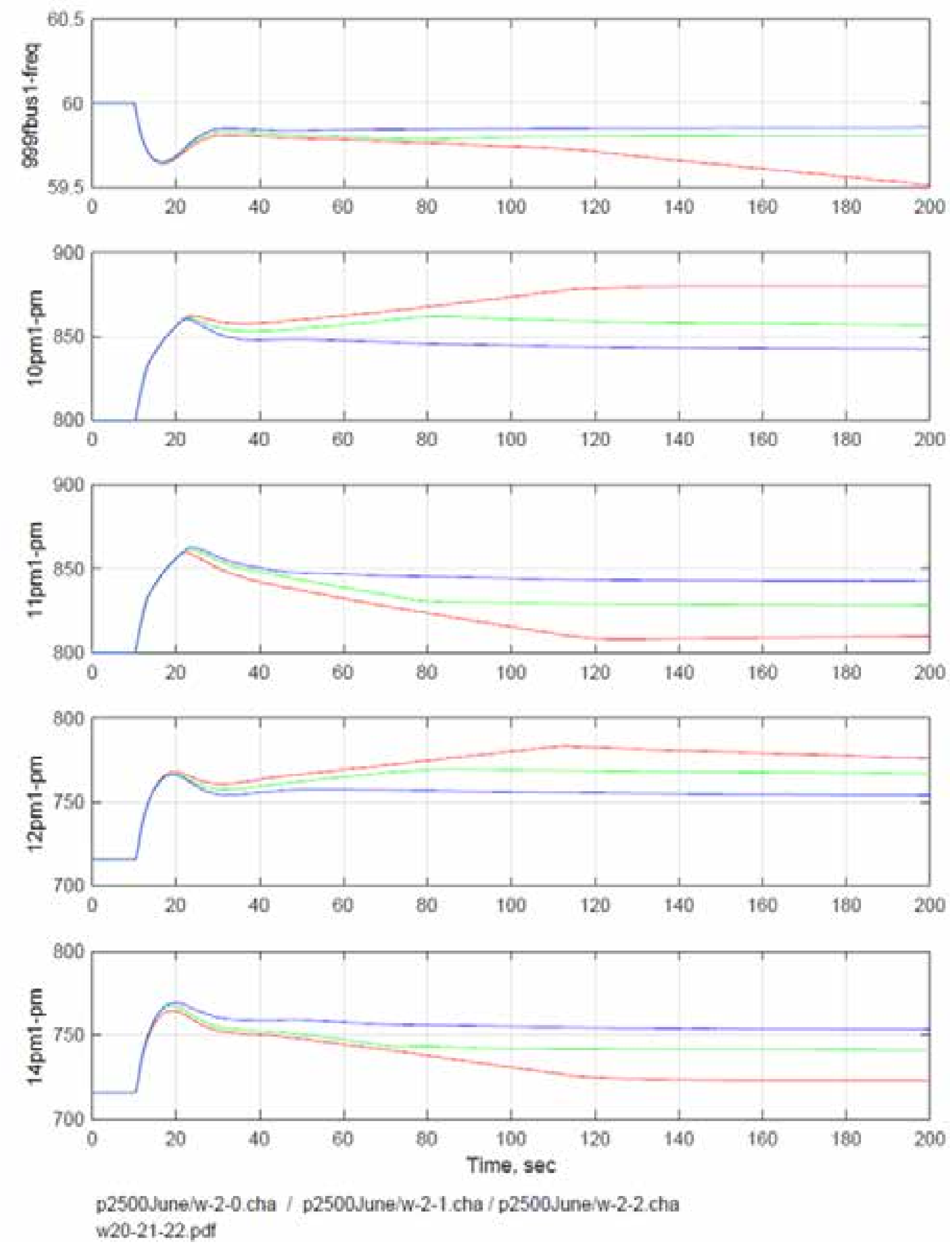

Figure 32. Response of turbine powers with power plant frequency bias

Red - Bf p = 0 - no frequency bias
Green - Bf p = 10 - frequency bias corresponds to half of governor droop
Blue - Bf p = 20 - frequency bias corresponds fully to governor droop

12. Deadband Effects

12.1 Programmed Deadband and Mechanical Imperfection

No machine is perfect; 'unintentional' deadbands due to mechanical realities and imperfections surely exist to varying extents in turbine controls However, the deadbands that are found in the mechanical actuators of modern digital electrohydraulic turbine controls are so small that they can be regarded as negligible in relation to deadbands that are placed deliberately in the logical parts of the controls. Accordingly, this discussion addresses deadband on the basis that deadbands of mechanical origin are negligibly small and that, for various reasons, deadbands are introduced 'intentionally' into the digital logic of modern turbine controls.

Figure 33 shows a selection of locations at which deadband logic can be introduced into the logic of proportional and proportional-integral turbine governors. This logic is typically embodied in digital code executed in 'programmed logic controllers' (PLCs).

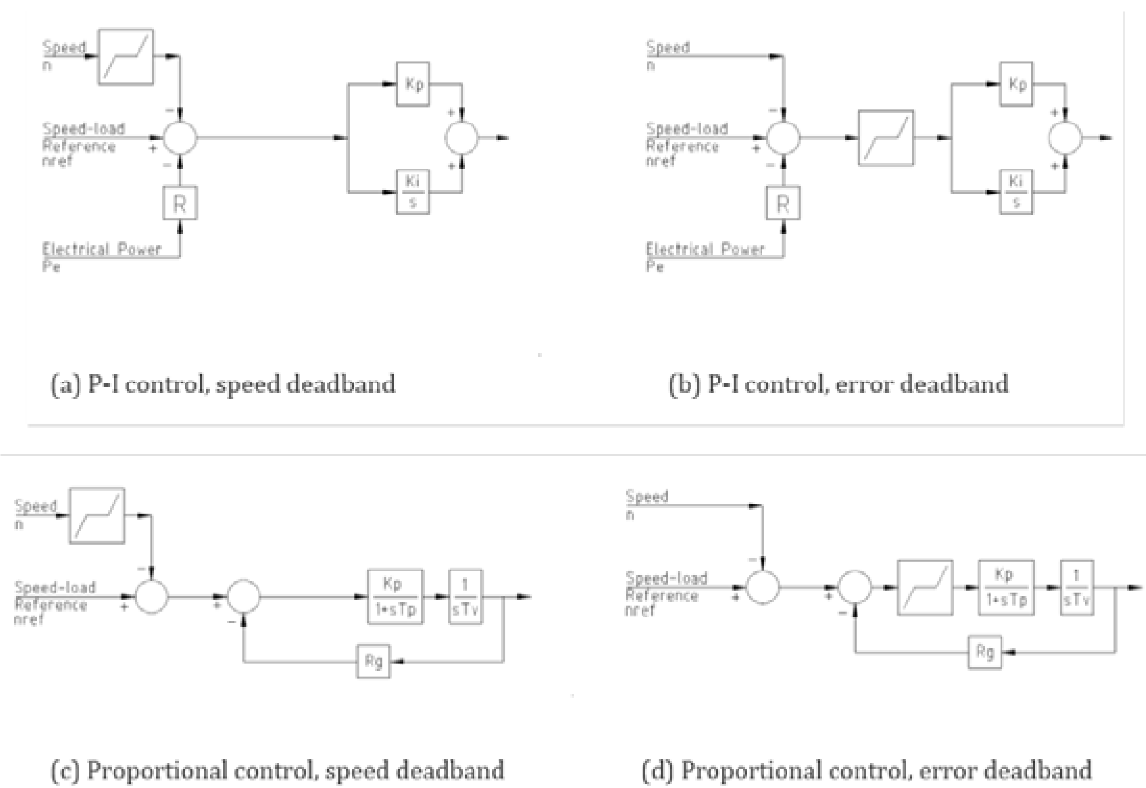

Figure 33. Locations of deadband in turbine controllers

12.2 Effect of Deadband on Turbine Behavior

Figure 34 and Figure 35 show recordings taken from a large single-shaft gas turbine running at substantially constant power output for 40 minutes. The green trace in Figure 34 shows that the governor is active even though speed variation is very small (+/- one RPM for most of the recording interval). The governor moves the control valve actively for much of the recording period but there are several intervals of as long as one minute in which the valves do not move. Figure 35 shows two plots of input versus the output of the governor's P-I control element and shows (within the limits of recording precision) that a deadband is present with a magnitude of 0.025 percent in terms of speed error. That is, a deadband in speed error terms of 0.00025*60 = 15 mHz.[6] It can be noted that this deadband is not due to imperfection in the machine, it is intentionally introduced in the digital code of the governor.

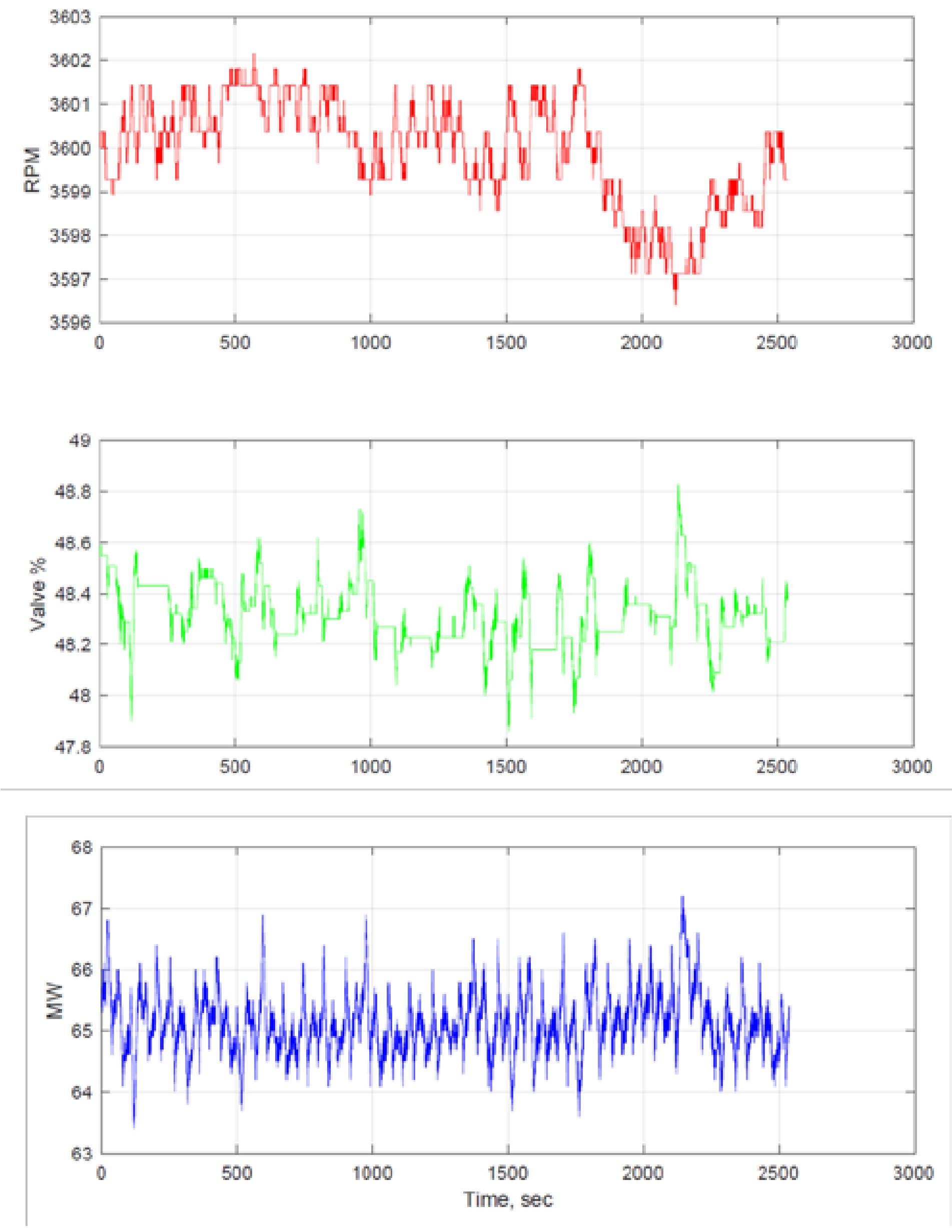

Figure 34. Behavior of single-shaft heavy-duty gas turbine running at near-constant part-load output

[6] For background information: The UK Grid Code (September 2016) requires that the deadband of gas turbine governors be not greater than +/-0.03 percent. The Italian Grid Code (July 2008) requires that governor deadband be not greater than +/- 0.02 percent.

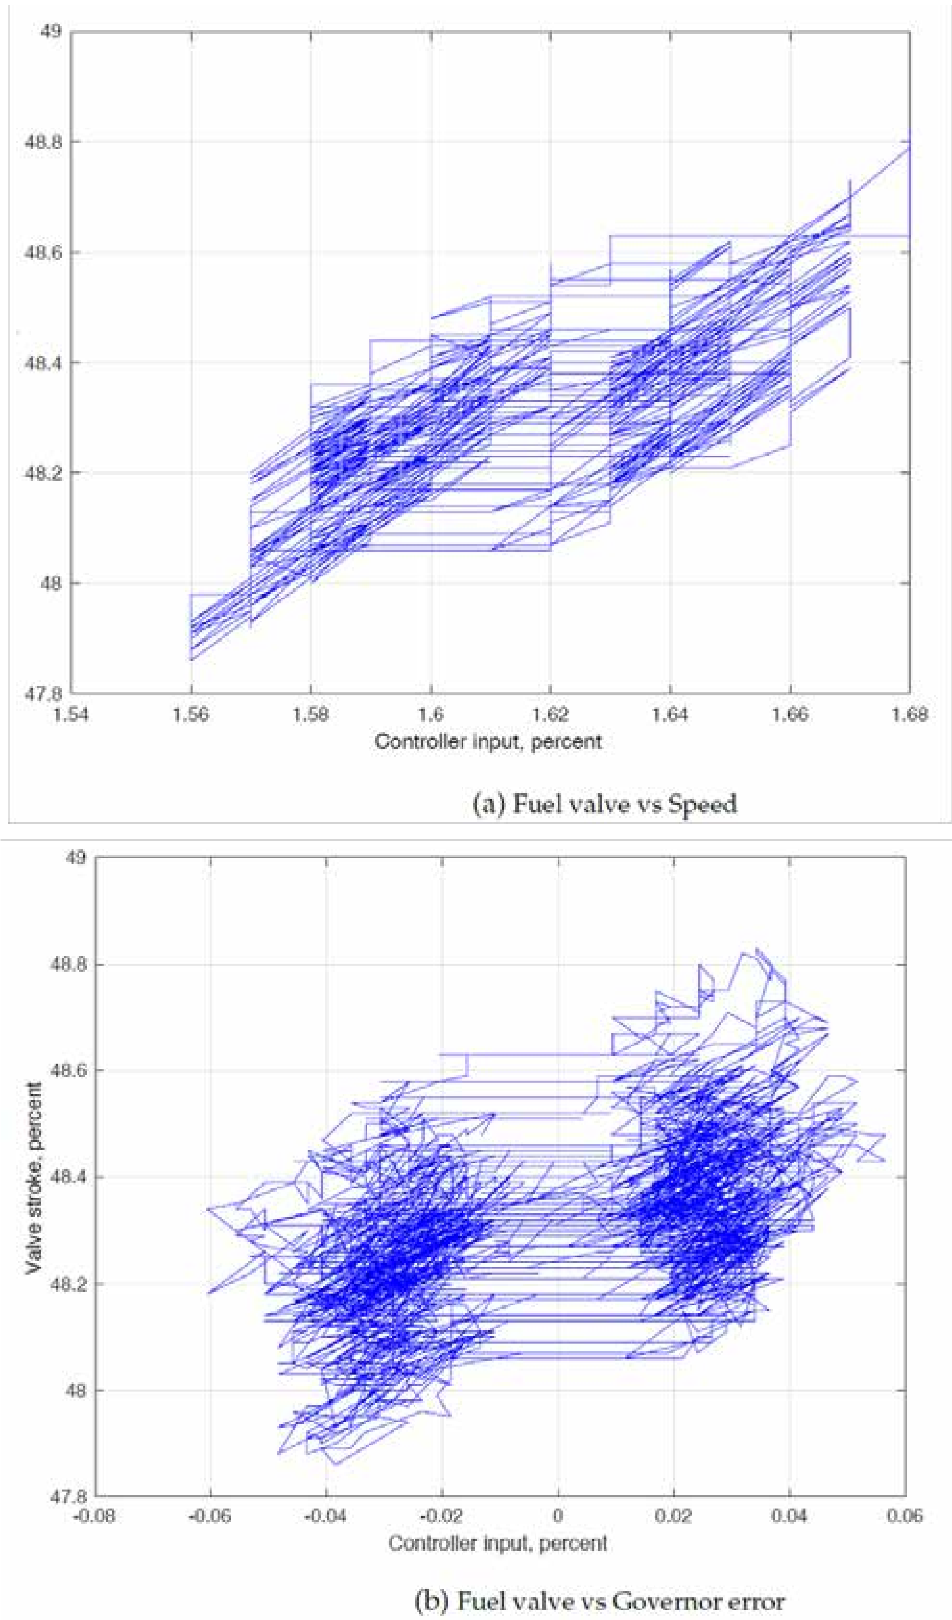

Figure 35. Governor deadband as seen in response of gas turbine control valves

12.3 Effect of Deadband on System-wide Frequency Response

12.3.1 Simulation model

To show the effect of governor deadbands on system wide frequency control behavior, the system model described in Section 6 has been altered as follows:

- The governors of the steam turbine units (SteamA and SteamB) have been changed from ggov1 to tgov1m in which deadband is located as shown in Figure 33(d)
- The deadbands of all thermal unit governors have been set to a common value (0.02 percent or 0.06 percent in most instances)
- Hydro units are unchanged from prior simulations and are assigned no deadband
- A random zero mean component has been added to the steady system load; a sample of the random load is shown in Figure 36

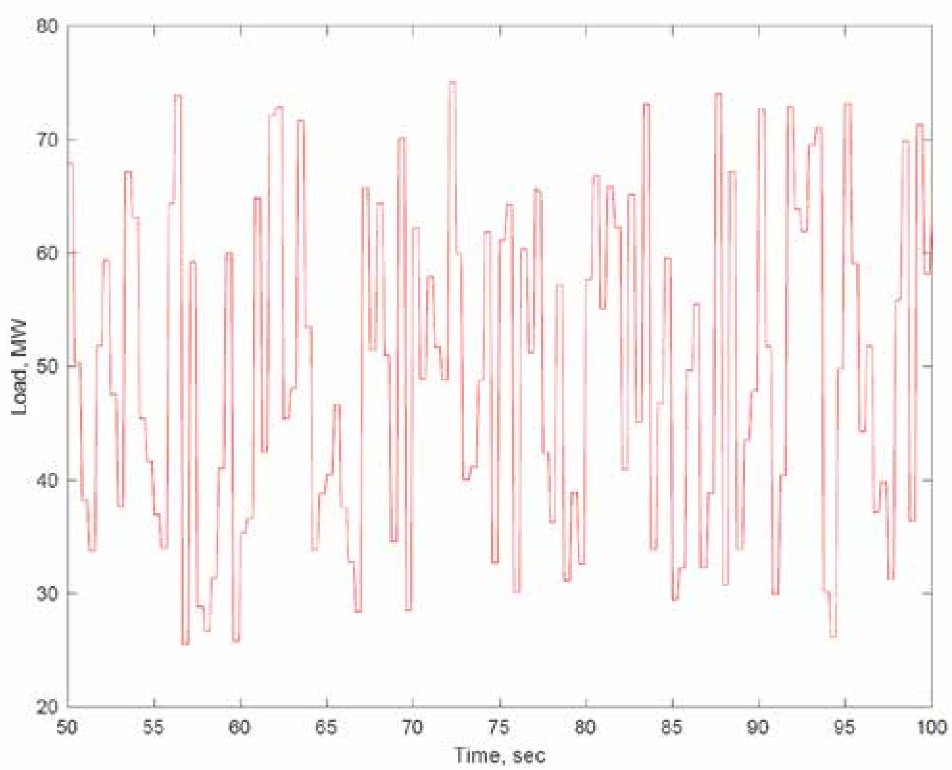

Figure 36. Form of random load component

12.3.2 Simulated system response

Table 10 summarizes simulations of the response of the system in the presence of governor deadband.

Table 10. Simulations including governor deadband effects

Case	Figure	Deadband (%)	Deadband (milliHz)	Power Step (%)
71	Figure 39	0.02	12	0
72	Figure 40	0.02	12	0.5
73	Figure 41	0.02	12	1.0
74	Figure 42	0.02	12	2.0
75	Figure 43	0.06	36	0
76	Figure 44	0.06	36	0.5
77	Figure 45	0.06	36	1.0
78	Figure 46	0.06	36	2.0

These simulations were made for the following composition of the generating fleet:

- responsive fraction, *Rfrac* = 0.4
- sustaining fraction, *Sfrac* = 0.9
- electronically coupled fraction, *Efrac* = 0.1

One half of the rotating generating fleet is not responsive to change of frequency and the fraction of the fleet that sustains its initial response is 0.36.

Figure 37 through Figure 44 show simulation results. In these figures red traces show simulated response with deadband and blue traces show response with no deadbands. The random variation of load is re-seeded for each simulation.

Start by comparing Figure 37 and Figure 41 which show response when there is no sudden power unbalance event other than the random variation of load. The blue traces show continual very small governor action when deadband is absent. The red traces show significant intervals with no motion of the valves when deadband is present. The red frequency traces in the two figures suggest that the larger deadband has allowed slightly larger random variation of frequency in response to the random variation of load. (Formal calculation of the statistical properties of the frequency variation has not been done.)

Next, compare Figure 44 and Figure 42 which show the response to sudden power unbalances of 2.0 percent and 0.5 percent respectively. As would be expected, the presence of deadband is strongly indicated (by valves remaining stationary) in the case of the small disturbance, but is essentially invisible when the disturbance is 2.0 percent of the system load.

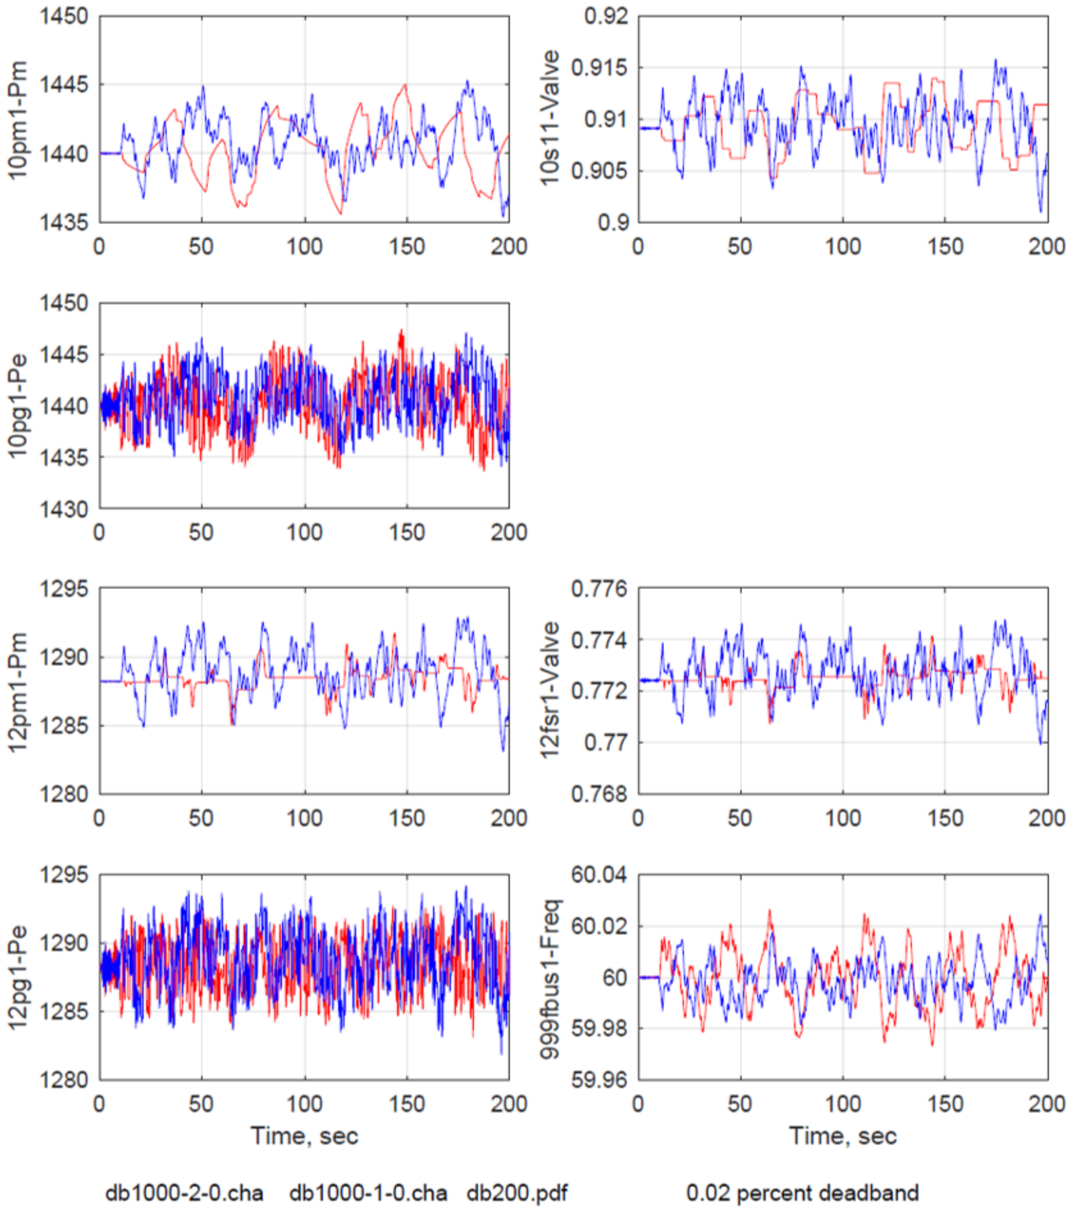

Figure 37. Simulation with deadband = 0.02%; Power step = 0.0%

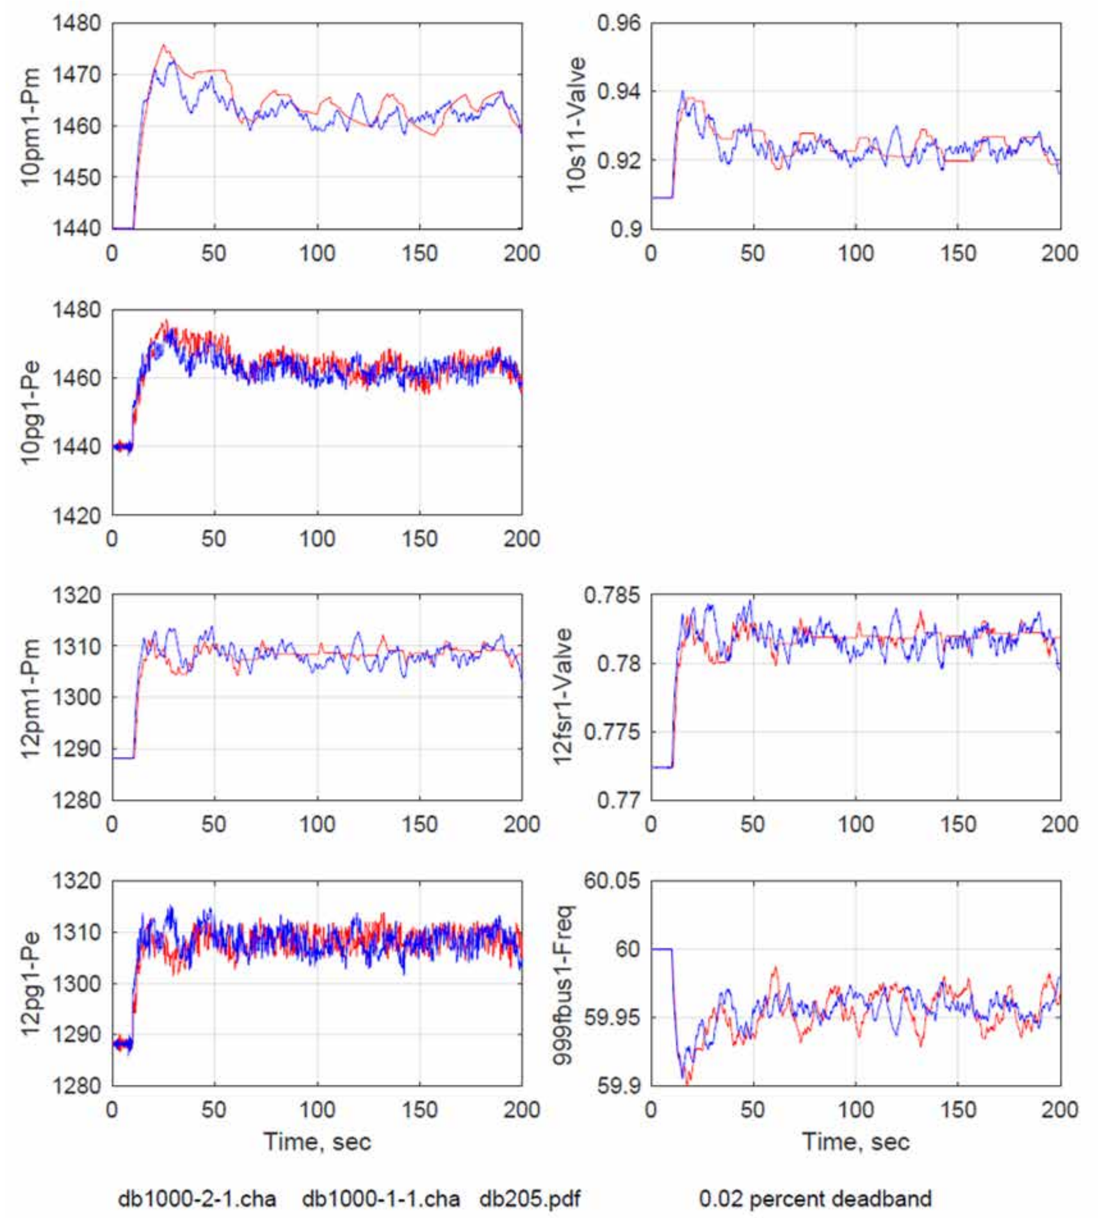

Figure 38. Simulation with deadband = 0.02%; Power step = 0.5%

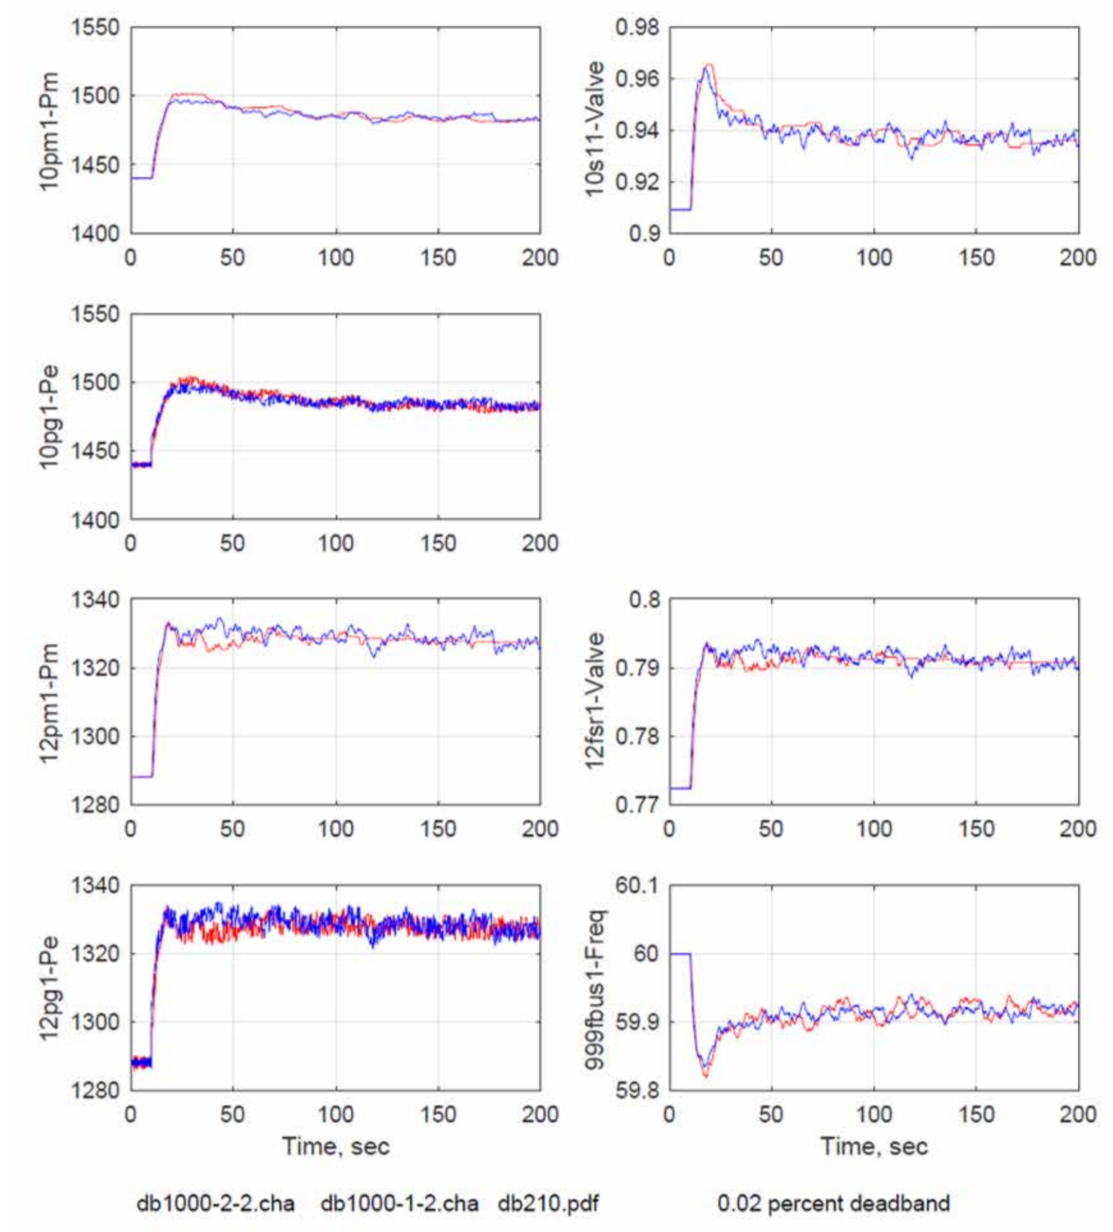

Figure 39. Simulation with deadband = 0.02%; Power step = 1.0%

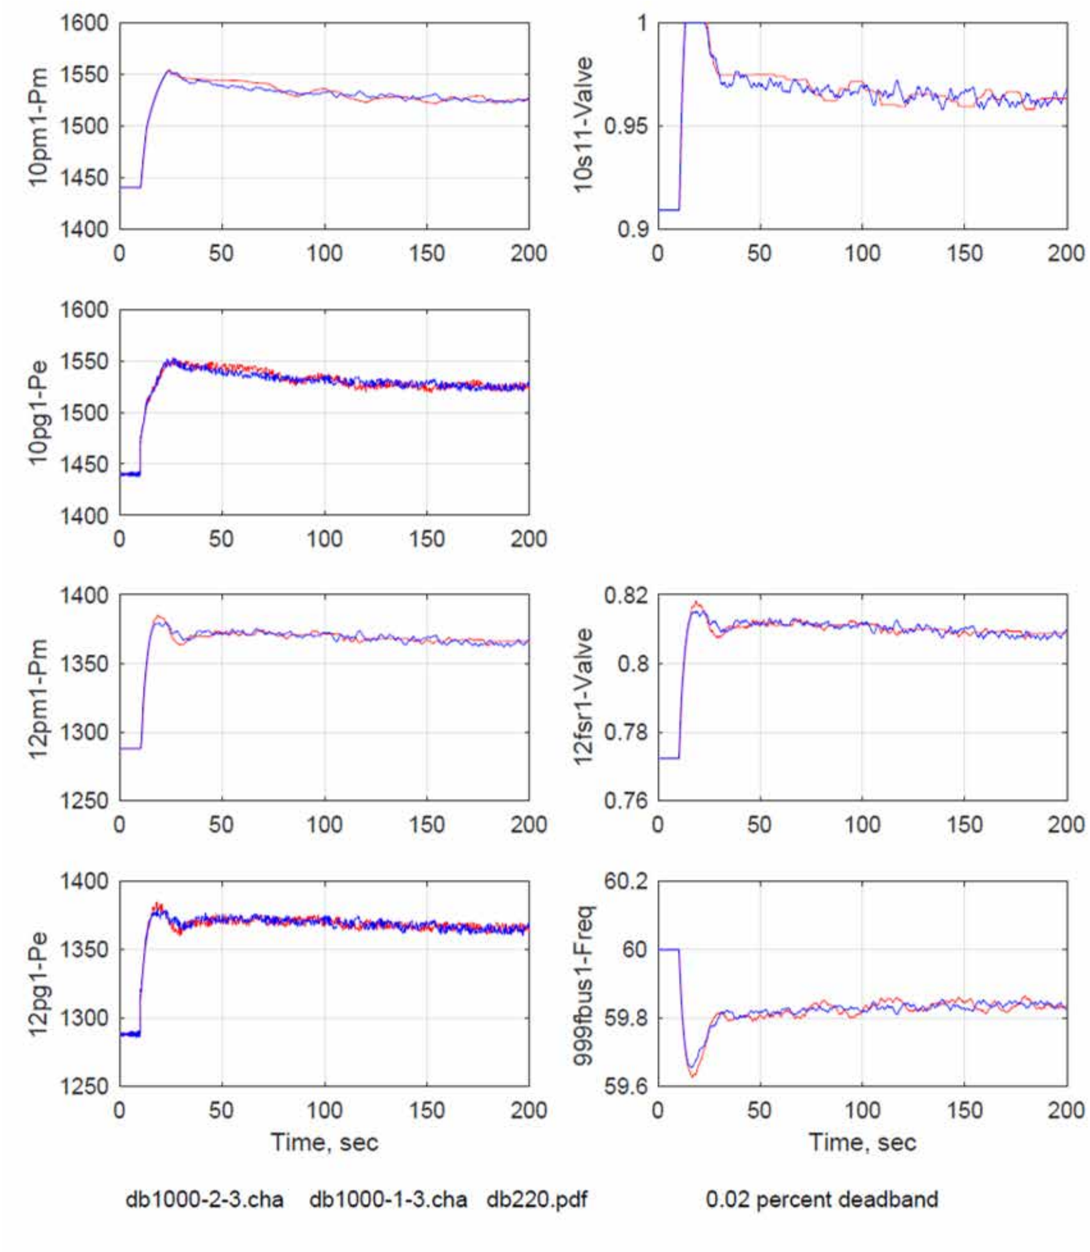

Figure 40. Simulation with deadband = 0.02%; Power step = 2.0%

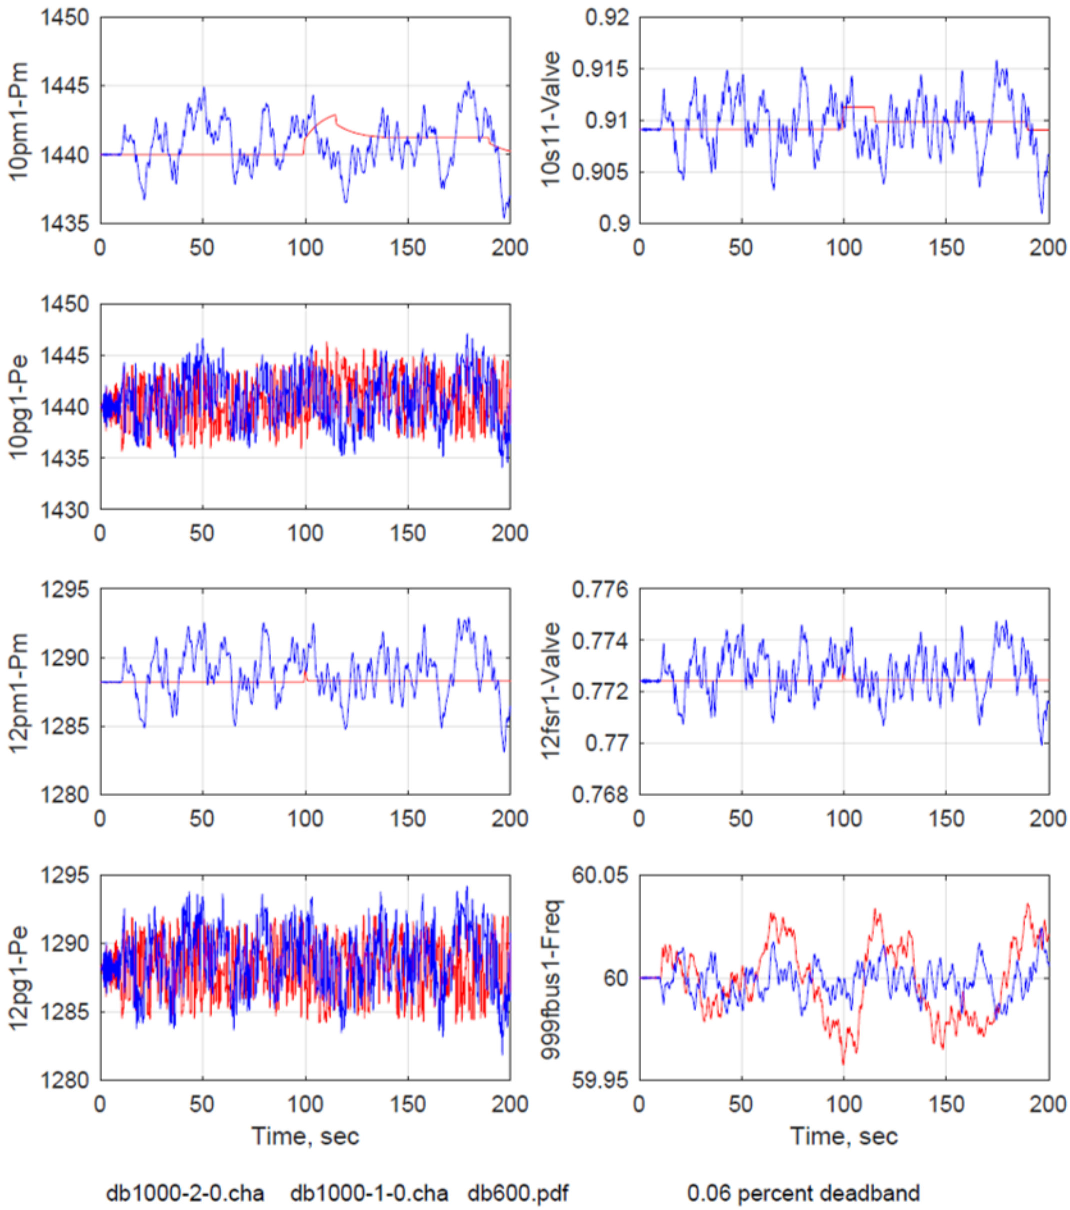

Figure 41. Simulation with deadband = 0.06%; Power step = 0.0%

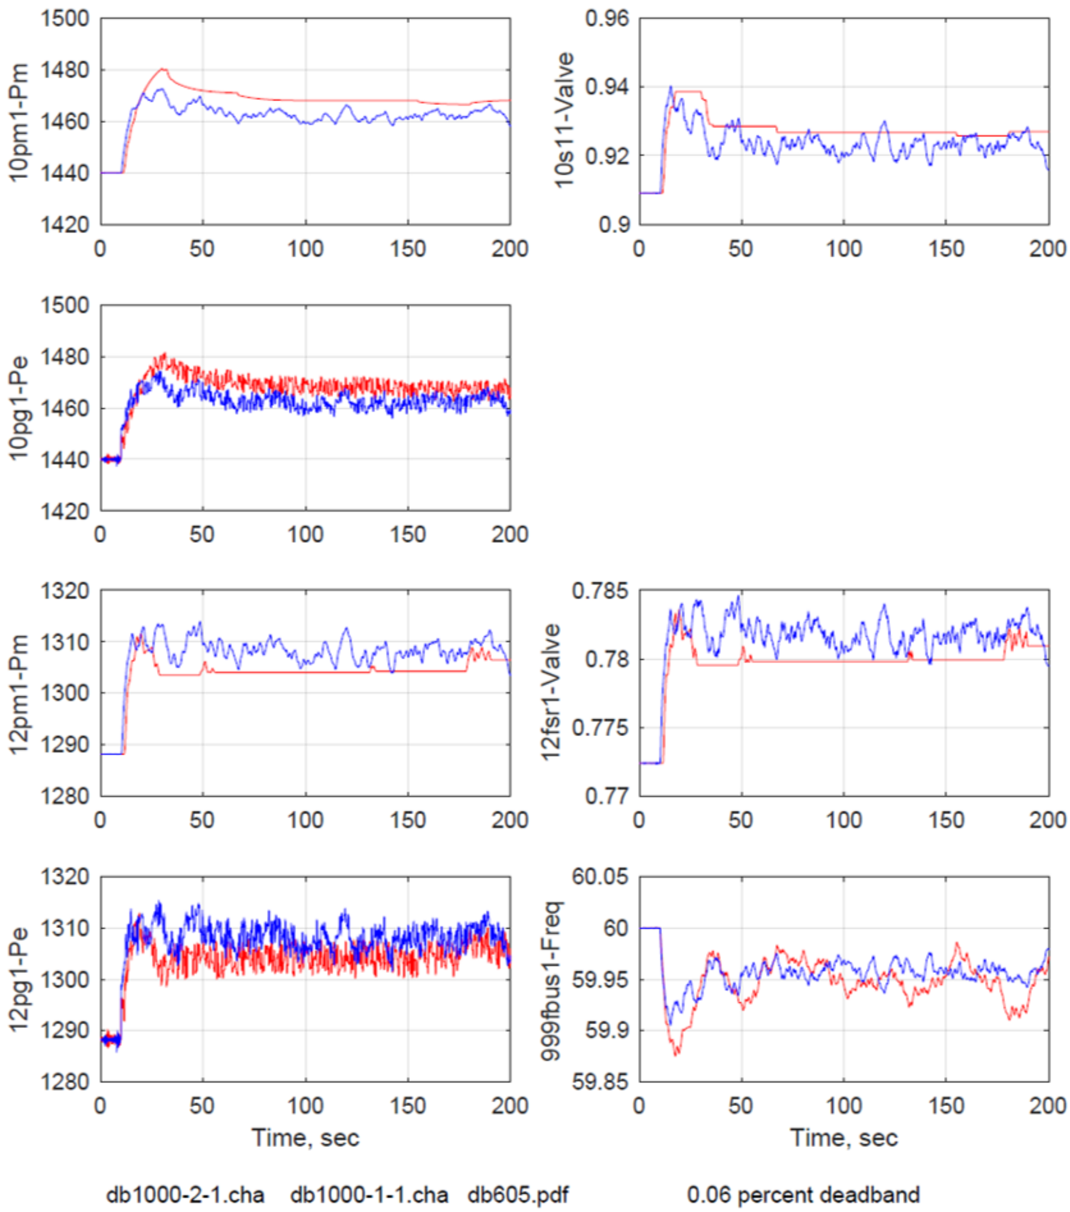

Figure 42. Simulation with deadband = 0.06%; Power step = 0.5%

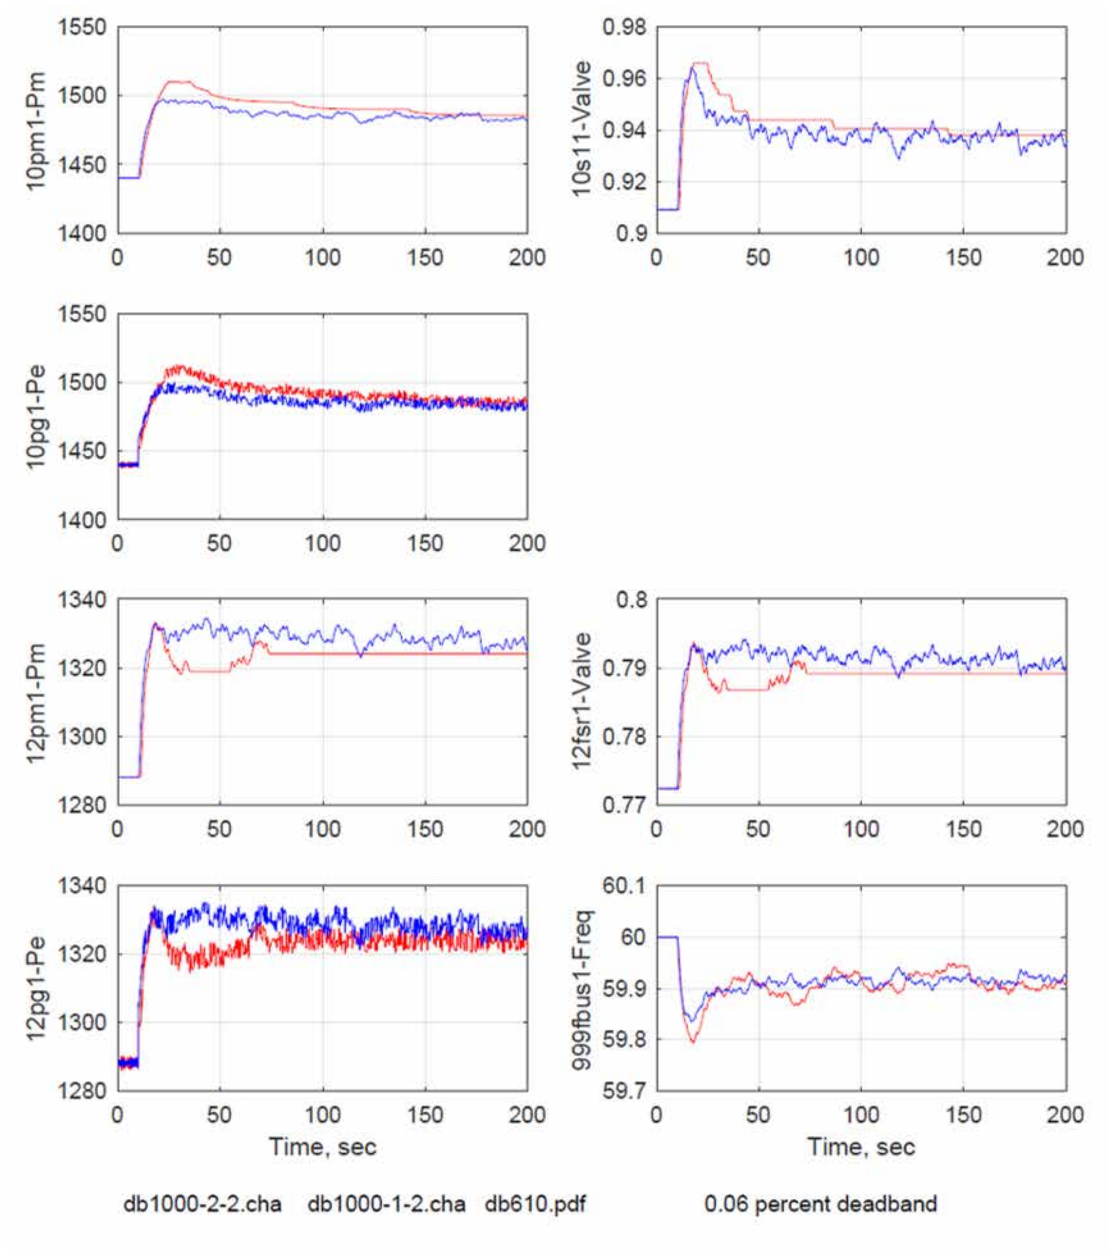

Figure 43. Simulation with deadband = 0.06%; Power step = 1.0%

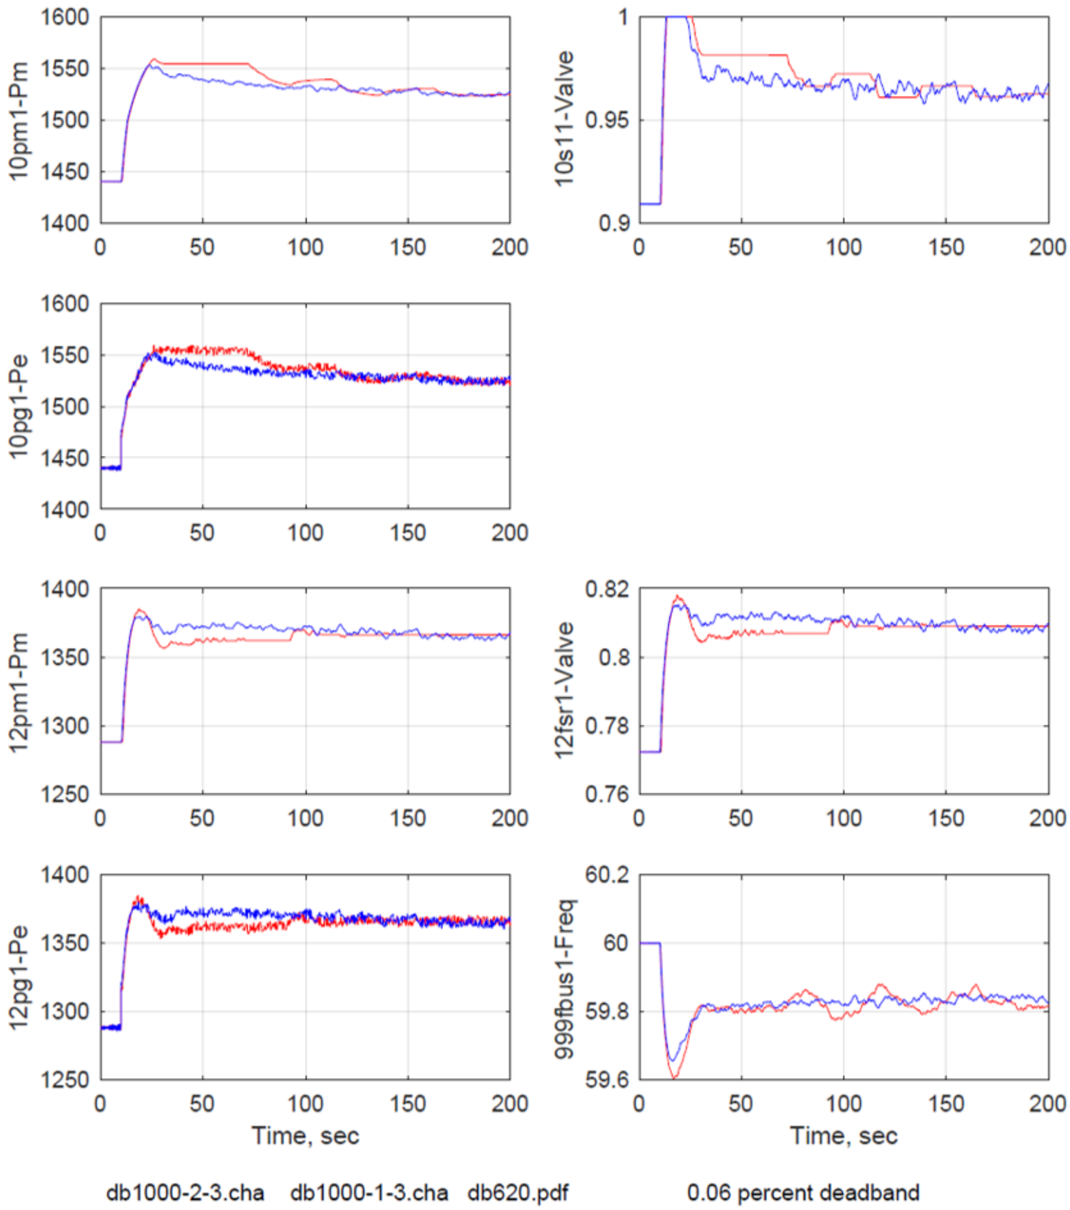

Figure 44. Simulation with deadband = 0.06%; Power step = 2.0%

Figure 45 summarizes the deadband simulations. The left hand two columns show simulations made with deadband set at 0.02 percent; the right hand columns were made with deadband set to 0.06 percent. The most apparent difference in system behavior is seen in the plots of frequency after the system has recovered from the sudden disturbances. It is apparent visually, though not by formal calculation, that the random variations of frequency are greater in the cases with 0.06 percent deadband than in the cases with small deadband.

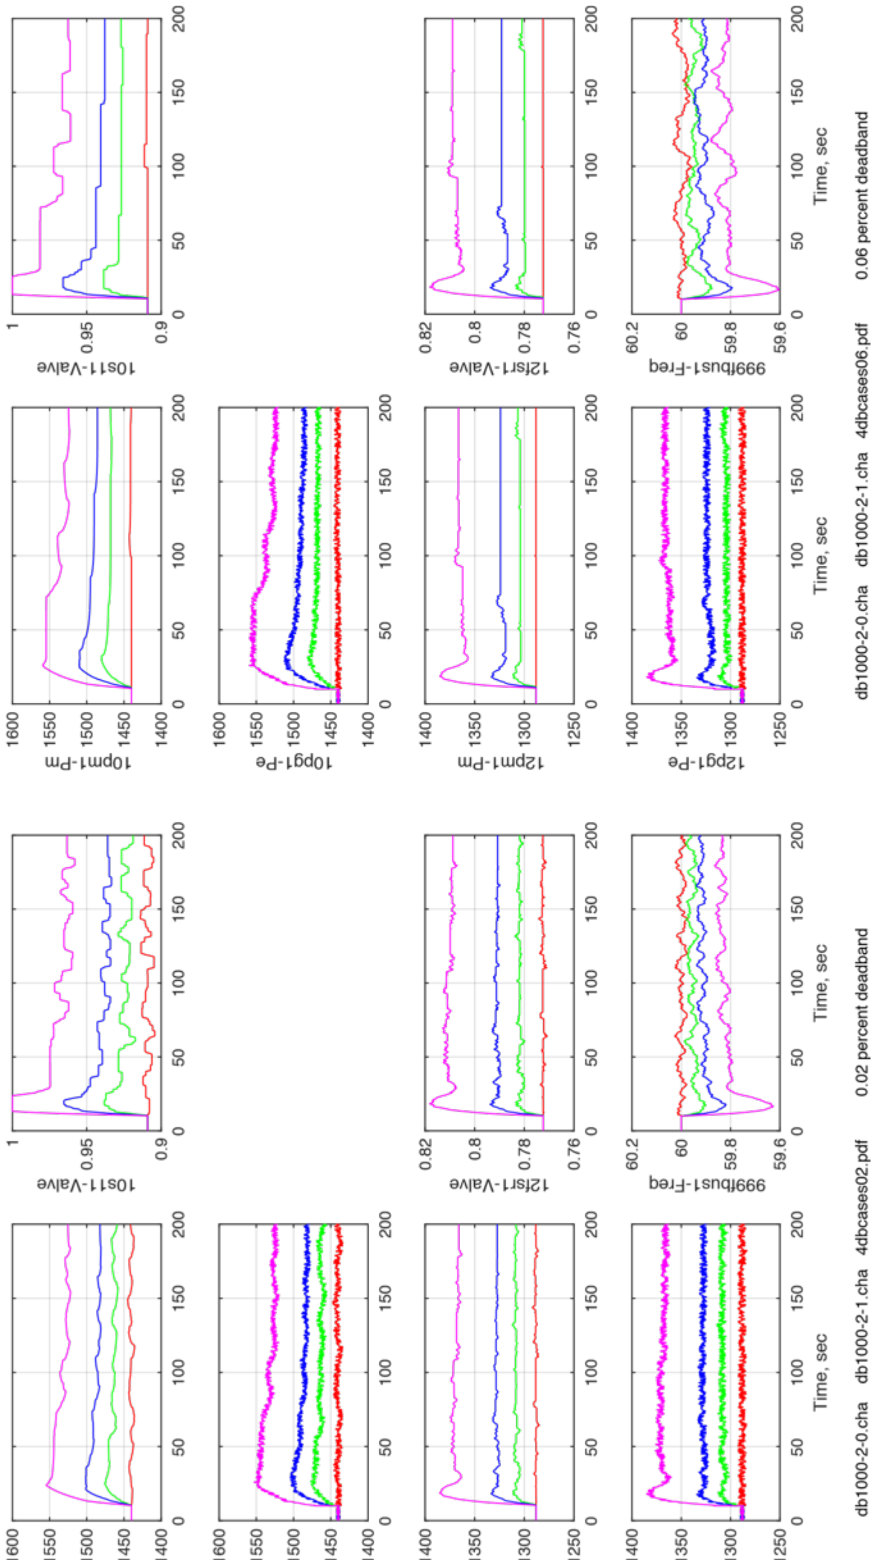

Figure 45. Simulation with deadband = 0.02 and 0.06%; Power steps = 0.0, 0.5, 1.0, 2.0%

13. Load Shedding as Primary Control Action

13.1 Programmed Load Shedding

This section considers the use of load shedding as a contributor to primary frequency response of the system, as distinct from a last-resort protective measure. Load shedding for this control-oriented purpose is typically distinguished by being activated at significantly higher frequencies than protective load shedding and by longer time delays before being activated. For the illustrative examples shown below a load shedding scheme (LR) is presumed to operate as follows:

- Load shedding to aid in frequency control is triggered at 59.7Hz and actuated with a definite time delay of 1 to 3 seconds
- The load circuit breaker operating time is 0.05 second
- The amount of load shedding triggered at 59.7Hz is 1, 2, or 3 percent of the total load
- Two additional protective load shedding blocks of 0.5 percent are triggered after 0.5 second delay at 59.3Hz and 59.0Hz

This load shedding scheme is similar to the LR scheme used in the ERCOT system, but is not intended to reproduce its operation in detail.

The intention of using load shedding as a primary control mechanism is to extend the size of generation deficit that the system can withstand without having its frequency decline to the levels where protective load shedding takes place. Here, for illustration, we consider that the system would operate with sufficient standard primary control margin to withstand a 2 percent generation deficit without any load shedding, but would rely on its LR scheme to deal with deficits greater than 2 percent.

13.2 Illustrative Simulations

13.2.1 Two and four percent generation deficits

Figure 46 and Figure 47 illustrate the role played by the LR scheme when the system faces generation deficits of 2 percent and 4 percent respectively.

The generation deficit of 2 percent is shown by Figure 46 to send frequency down an initial nadir at just above 59.7Hz. The arrest of frequency decline cannot be sustained as secondary controls withdraw production and the resulting slow decline after the initial arrest results in frequency reaching 59.7Hz at 26 seconds; the LR load shedding is triggered. The red, green, and blue traces show the frequency recovery following the disconnection of 1, 2, and 3 percent of the total load. The shedding of 2 percent (green) is the ideal in that it allows the generating fleet (less the unit that tripped) to return to the initial power output at nominal frequency. Shedding 1 percent allows frequency to recover but only as far as 59.9Hz. Shedding 3 percent of the load results in an overfrequency. The below nominal and above nominal settling frequencies would be corrected in due course by system-level secondary control actions.

Figure 47 corresponds to Figure 46 but shows the response when the generation deficit is 4 percent. Frequency goes down to 59.7Hz very quickly; there is no possibility that primary response of rotating generation will arrest the decline and the LR scheme operates promptly. The red trace shows that the shedding of 1 percent of the total load is just sufficient to arrest the rapid decline but not enough to maintain the arrest as secondary controls withdraw power. Frequency resumes its (slower) decline and the decline is finally stopped by the operation of the third stage of LR action at 59.0Hz. The green and blue curves show that the shedding of 2 or 3 percent of the load, combined with the action of the responsive-sustaining part of the rotating fleet, is enough for a sustained recovery of frequency.

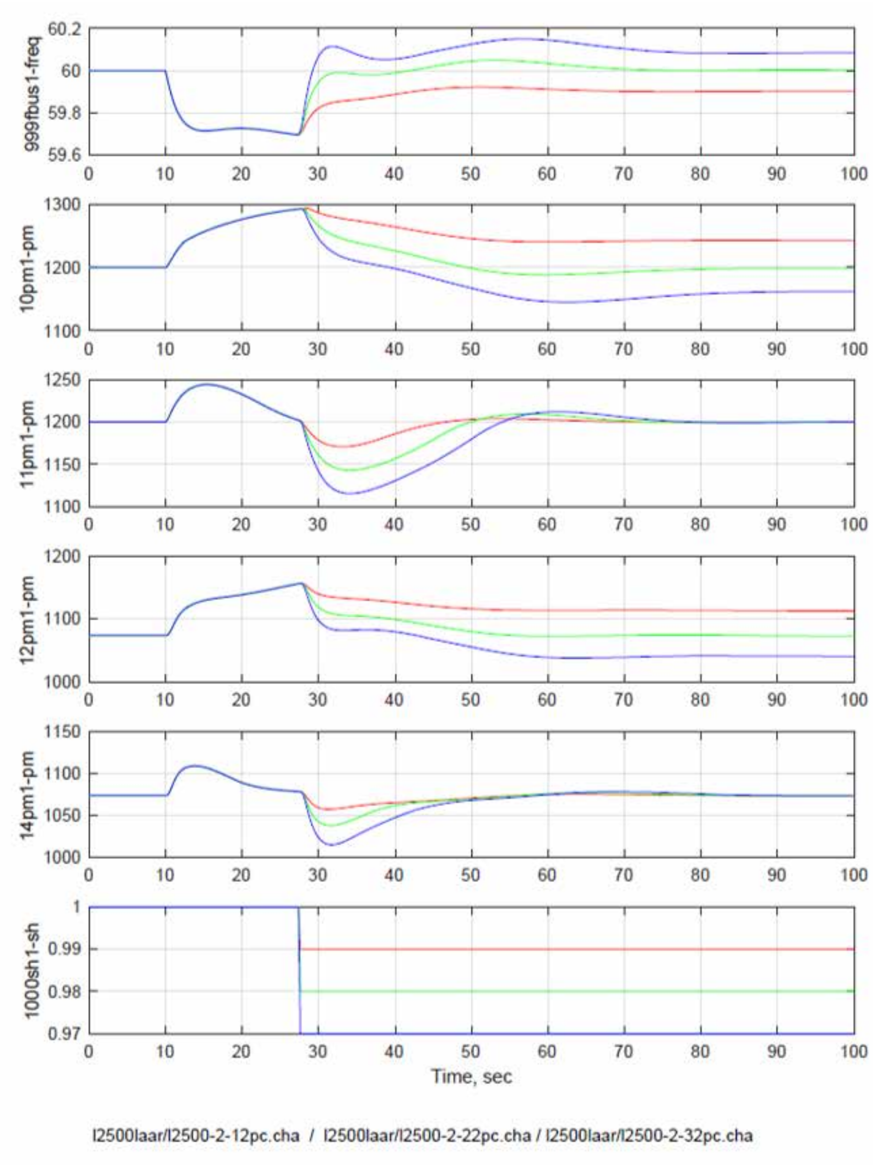

Figure 46. Response of system to 2% generation deficit

Red - LR sheds 1%
Green - LR sheds 2%
Blue - LR sheds 3%
LR triggered at 59.7Hz with delay=1.0 second

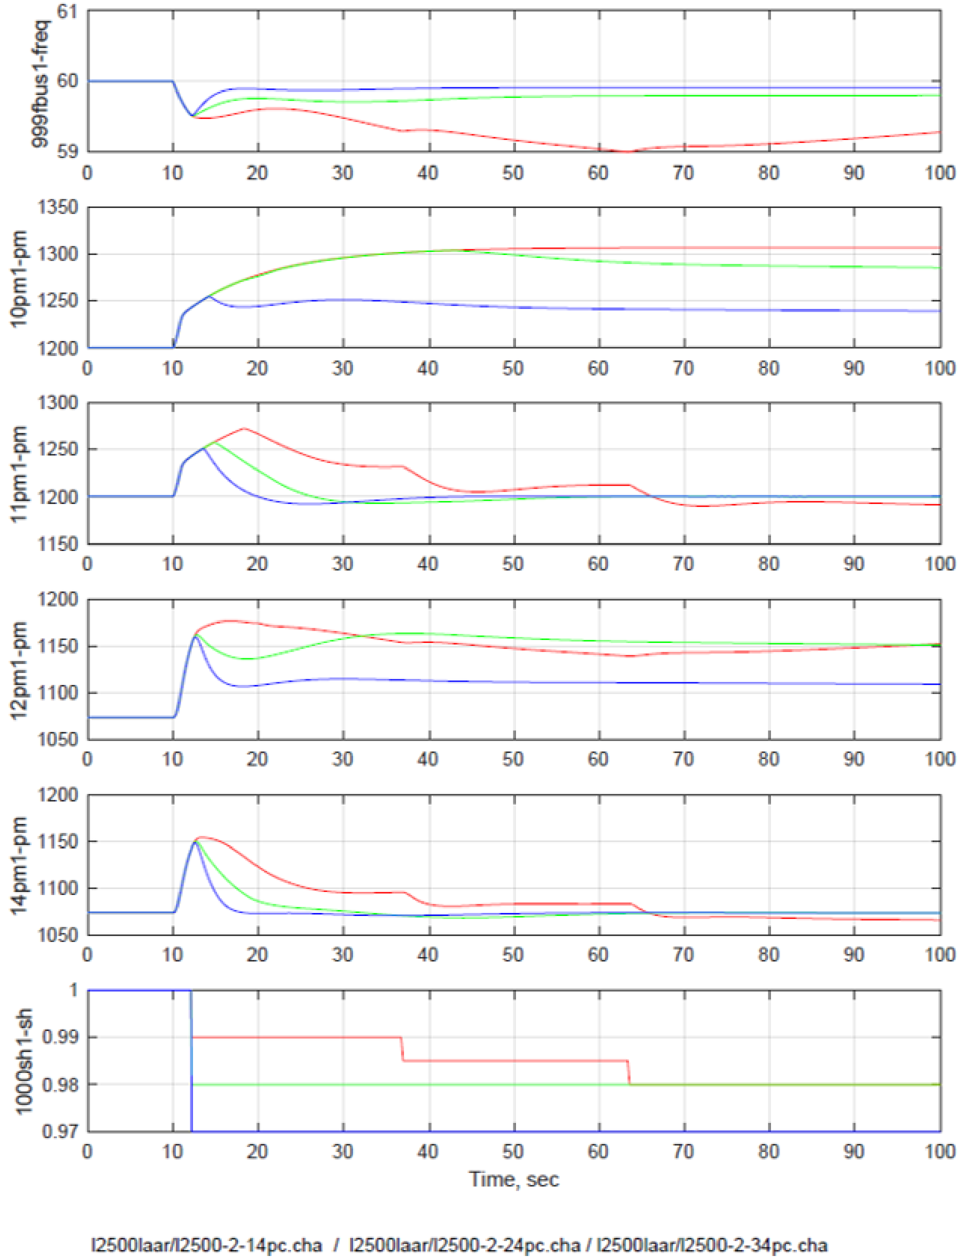

Figure 47. Response of system to 4% generation deficit

Red - LR sheds 1%
Green - LR sheds 2%
Blue - LR sheds 3%
LR triggered at 59.7Hz with delay=1.0 second

13.2.2 Relationship to inertia constant

Figure 48 shows simulations of response to a two percent generation deficit with the rotating fleet assigned inertia constants of 2, 4, and 6 seconds. The significant result of this comparison that the behavior of the system is very little changed by the very broad change of inertia.

The primary response of the rotating-responsive fleet is just able to arrest the initial frequency decline. Changing the inertia constant changes the time at which arrest occurs but does not significantly change the frequency. In the three cases the initial arrest is at a frequency just above the LR triggering frequency of 59.7 Hz. There is not enough sustaining-responsive generation and frequency slides downward after the initial arrest. LR is triggered at about 26 seconds and, by shedding of 1 percent of the load, saves the system.

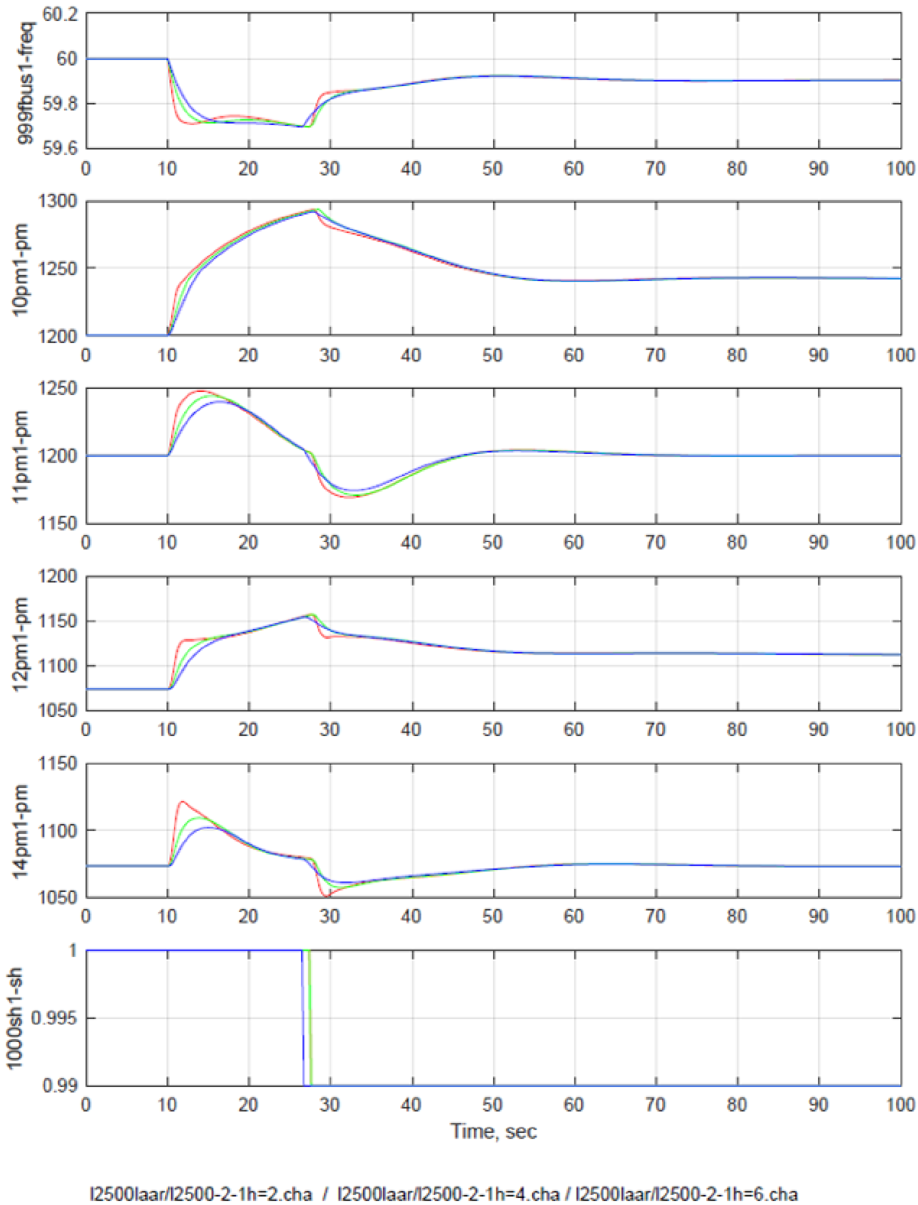

Figure 48. Variation of system behavior with variation of inertia constant

Generation deficit = 2%
LR Sheds 1% when triggered at 59.7Hz
Red - H=2
Green - H=4
Blue - H=6

14. Load Frequency Sensitivity

14.1 Facts and Assumptions

The characteristic of the power system load with regard to frequency has generally been regarded as being favorable to the control of the system. That is, it has generally been assumed that the total system load decreases as frequency decreases. A common assumption regarding load is that it behaves essentially as stated by Equation 4 in Section 6.5:

$$P_{load}(f) = P_{load}(f_{nom}) * (1 + D_l\,(f - f_{nom}))$$

Equation 6

Historically, the makeup of the load provided a degree of guidance as to the value of the load damping factor, D_l. When load consisted predominantly of incandescent lighting and resistance electric heating, the power consumed was substantially independent of frequency and it could be assumed with some confidence that $D_l = 0$. When the load was known to be predominantly motors driving fans and pumps, whose power consumption rises as speed is increased, it was reasonable to assume that the load damping factor was in the range of one to two.

The historical guidance regarding the value of D_l is fading fast, however, as electronically controlled load takes over as the predominant component of system load. As electronic controls are interposed between end loads, (lamps, computers, fans, pumps, etc.) the programs of the electronic controllers are becoming the dominant influence on the characteristic of the load that the generating resources must meet. There is reason for concern that this trend is making the load less favorable with regard to the control of frequency. A key factor in this is the rather natural tendency of users and equipment manufacturers to expect a motor controller to run its driven load, (perhaps a conveyer belt in a factory, perhaps the milking machine on a dairy farm) at a speed chosen by the user, as distinct from a speed that is advantageous to the grid. A speed-sensitive load, such as a pump, whose electronic controller constantly adjusts itself internally to maintain constant speed presents a frequency-independent load to the grid and thence to the generating fleet. The trend can be assumed, therefore, to be for the value of D_l to be decreasing which is in the direction that is detrimental to the control of grid frequency.

14.2 Effect of Load Damping

Figure 49 shows the sensitivity of the frequency trajectory to the load characteristic for a two percent generation loss when 60 percent of the generating fleet is responsive to frequency, *Rfrac* = 0.6, and a half of that responsive generation sustains its contribution, Sfrac = 0.5. The system inertia constant is four. It is clear that the frequency sensitivity of the load contributes measurably to both the initial arrest of frequency decline and the subsequent recovery.

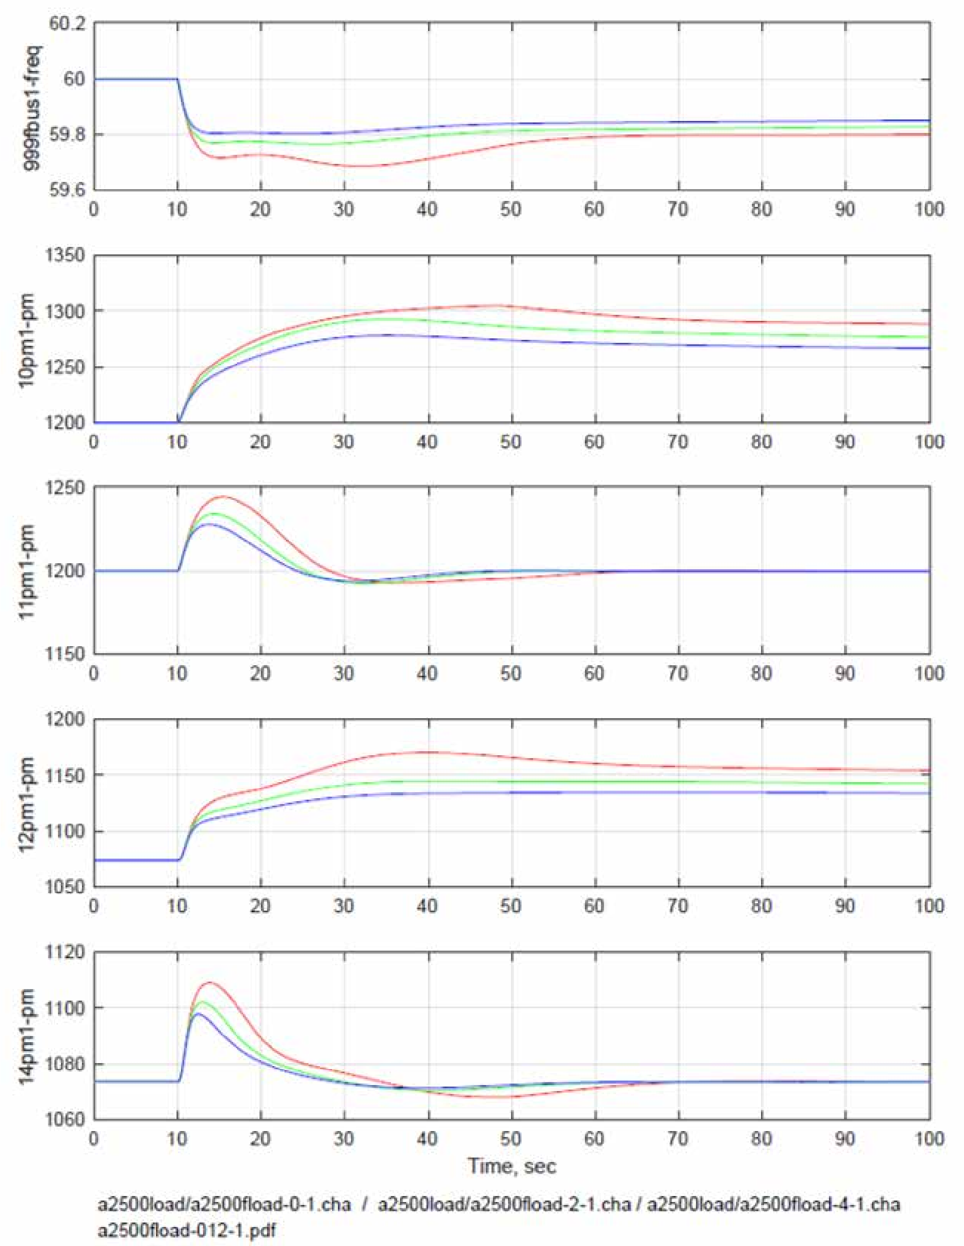

Figure 49. Effect of load frequency sensitivity on system frequency

Generation deficit = 2%
Red - D=0
Green - D=1
Blue - D=2

14.3 Load Damping in Relation to System Inertia

Figure 50 compares the effect of the load frequency characteristic with the effect of changing system inertia. The red and green curves in Figure 50 show simulations made with the system-wide inertia constant set to H = 3 and with the load damping factor set to D_l = 0 and D_l = 2, respectively; the effect of changing the load characteristic, D_l, is very clear.

Next, look at the red and blue curves in Figure 50. These two curves were made with the same load characteristic, $D_l = 0$ but with the inertia constant set to $H = 3$ and $H = 4$, respectively. The difference in the frequency trajectories when inertia is changed by 33 percent is barely visible.

These figures are a quite strong indication that the changes that are presently occurring in the *load* side of the industry are as significant as those associated with the generating fleet.

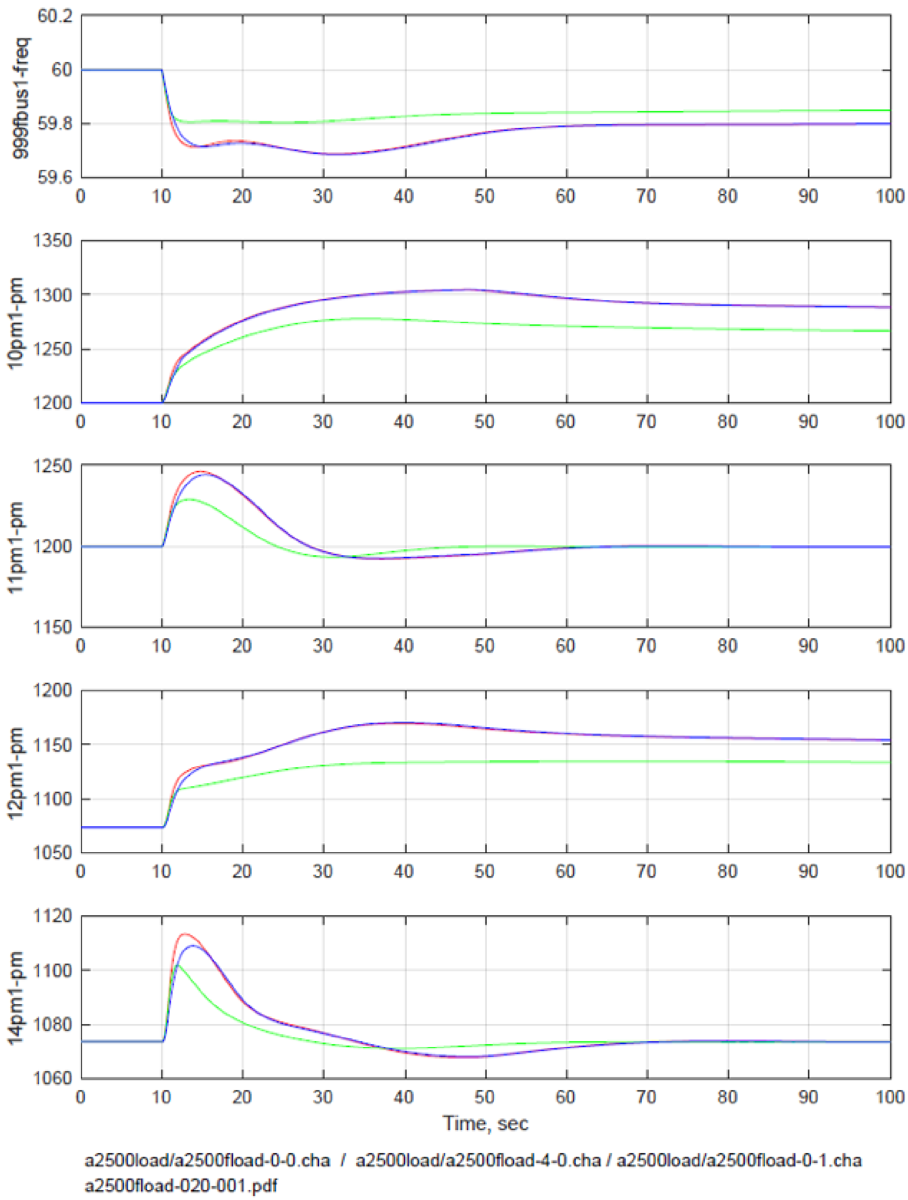

Figure 50. Relative effect of load frequency sensitivity and inertia constant on system frequency

Generation deficit = 2%
Red - D=0, H=3
Green - D=2, H=3
Blue - D-0, H=4

15. References

Maxwell, J.C. (1867). "On Governors." Proceedings of the Royal Society of London. January 1867. doi:10.1098/rspl.1867.0055.

Undrill, J.M. (2010). *Power and Frequency Control as it Relates to Wind-Powered Generation*. Lawrence Berkeley Laboratory. LBNL-4143E. December 2010. https://certs.lbl.gov/publications/power-and-frequency-control-it