Final Report:

Technology Capabilities for Fast Frequency Response

Prepared for: **Australian Energy Market Operator**

Prepared by: **Nicholas Miller, project lead**

Debra Lew

Richard Piwko

Contributions from:

Lou Hannett

Sebastian Achilles

Jason MacDowell

Matt Richwine

Douglas Wilson

Mark Adamiak

March 9, 2017

Imagination at Work

Legal Notices

This report was prepared by General Electric International, Inc. as an account of work sponsored by AEMO. Neither AEMO, General Electric International, Inc., nor any person acting on their behalf:

1. Makes any warranty or representation, expressed or implied, with respect to the use of any information contained in this report, or that the use of any information, apparatus, method, or process disclosed in the report may not infringe privately owned rights.

2. Assumes any liabilities with respect to the use of or for damage resulting from the use of any information, apparatus, method, or process disclosed in this report.

Foreword

This paper was prepared by General Electric International, Inc. acting through its Energy Consulting group based in Schenectady, NY, and submitted to AEMO. Questions and any correspondence concerning this document should be referred to:

Nicholas W Miller

Senior Technical Director

GE Energy Consulting

1 River Road

Schenectady, NY 12345

+1-518-385-9865

Nicholas.miller@ge.com

AEMO Statement

"This report was commissioned by AEMO in June 2016 under its Future Power System Security Program. The report forms part of the broad analysis AEMO is undertaking to identify and specify potential technical solutions to current and future power system security challenges."

SYNOPSIS

This report discusses Fast Frequency Response (FFR) as a potential mitigation option for the Australian Energy Market Operator (AEMO) in maintaining system security during low inertia conditions on their system. The report examines (1) the practicality, requirements and capabilities of several non-synchronous resources to provide FFR, (2) AEMO system requirements and benefits of FFR, and (3) considerations for procuring FFR. The key study findings are:

- The following “unconventional” technologies may be suitable for providing FFR: wind turbines; lithium, flow and lead-acid batteries; flywheel energy storage systems (inverter-interfaced and non-inverter systems); supercapacitor energy storage systems; solar photovoltaics (PV); load based resources; and high-voltage DC (HVDC) transmission. Inertia-based FFR (IBFFR, also known as “synthetic inertia”) from wind turbines can make a valuable contribution. PV could be an important FFR contributor in future. Wind and PV can also contribute to primary frequency response. Both synchronous and asynchronous flywheels have beneficial attributes that should be considered.
- The primary method to manage South Australia islanding should be by SPS (Special Protection System). An SPS should also be designed to address related risks, including loss of generation events (like the September 28, 2016 system black event). These low inertia conditions chart new territory with respect to system dynamics and stability, especially in South Australia. They should be addressed together with frequency response.
- Under low inertia conditions, identification of events must be quick and accurate. FFR response time is critical, but a balance between making high fidelity decisions to act and speed is needed. The analysis suggests that total response times on the order of one quarter to one half second are sufficiently fast.
- Detailed simulations are critical for determining the proper amount of FFR, and system needs and operational practice for that FFR. As inertia decreases, it becomes more important to provide the proper amount of FFR.
- Load tripping or control will continue to be part of risk management for South Australia and methods for this will need to evolve from present practice. Conventional under-frequency load shedding (UFLS) should not be relied on for extremely high rate-of-change-of-frequency (RoCoF).
- It is important to ensure system security first, and base market constructs on that foundation. Technologies with different energy delivery behaviors can be compared for efficacy. Costs of FFR should be compared to the variable cost of alternative operating strategies. Capital versus variable costs of FFR should be reflected in procurement. Some services, like IBFFR, should probably be handled separately.
- Achieving possible future conditions of zero synchronous generation, will require significantly more infrastructure, and evolution of technology and understanding.

The key recommendations going forward are:

- A comprehensive effort to specify and design a SPS to address separation of South Australia from the rest of the National Electricity Market (NEM) should be a high priority. The necessary system analysis and economic analysis should be performed to establish the right mix of new FFR technologies and direct load tripping for the SPS.
- Properly structured production simulations, that can capture these new constraints and consider provision of FFR and FCAS from a full range of technologies should be run. Economic evaluations, considering both capital and operating costs, are needed to provide meaningful comparisons between options. FFR options should include the technologies outlined in this report and conventional resources, as well as altered, more conservative operating strategies and relaxed performance standards.
- Model validation for evaluation of the dynamic performance of the NEM is needed, with attention to whether the available models provide adequate fidelity for simulations of the more extreme conditions under consideration. Performance of all important elements, including renewable generation and all resources providing FCAS must be considered.
- Detailed dynamic analysis of the NEM is needed. Illustrative cases and relationships developed with the very simple model of the South Australia system in this report provide guidance for relationships to be more formally calculated.
- For the near term, consider system solutions that will maintain a minimum level of inertia and that will contribute other performance benefits as well (i.e. improved transient, frequency and voltage stability).
- The possibility for utility scale solar PV to provide low marginal cost FFR by taking advantage of emerging practice of relative short-term ratings of panels and inverters should be investigated (this is an industry-wide issue and recommendation).

Table of Contents

Table of Figures

List of Acronyms

AC	Alternating Current
AEMO	Australian Energy Market Operator
AGC	Automatic Generation Control
AS	Ancillary Services
BA	Balancing Authority
BESS	Battery Energy Storage System
BMS	Battery Management System
BOP	Balance of Plant
CIGRE	International Council on Large Electric Systems
CRS	Contingency Reserve Service
CWFT	Continuous Wave Frequency Transducer
DC	Direct current
DER	Distributed Energy Resources
DFG	Double-fed Generator
DOD	Depth of Discharge
DOE	Department of Energy (U.S.)
DSM	Demand Side Management
EMT(P)	ElectroMagnetic Transients (Program)
EPRI	Electric Power Research Institute
ERCOT	Electric Reliability Council of Texas
FCAS	Frequency Control Ancillary Services
FESS	Flywheel Energy Storage System
FFR	Fast Frequency Response
FRM	Frequency Response Measure
FRO	Frequency Response Obligation
GE	General Electric International, Inc.

HELCO	Hawaii Electric Light Company
HVDC	High Voltage Direct Current
IEC	International Electrotechnical Commission
IESO	Independent Electricity System Operator (Ontario)
LCC	Line commutated converter
MPPT	Maximum Power Point Tracking (for photovoltaics)
MW	Megawatt
MWH	Megawatt hour
NEM	National Electricity Market
NERC	North American Electric Reliability Corporation
NYISO	New York Independent System Operator
OEM	Original Equipment Manufacturer
PCS	Power Conversion System
PFR	Primary Frequency Response
PLL	Phase-locked loop
PMU	Phasor Measurement Unit
PSH	Pumped Storage Hydropower
PSS	Power System Stabilizers
PWM	Pulse Width Modulation
PV	Photovoltaics
RAS	Remedial Action Scheme
RFP	Request for Proposal
RoCoF	Rate of Change of Frequency
RPS	Renewables Portfolio Standard
SA	South Australia
SCADA	System Control and Data Acquisition
SCE	Southern California Edison
SCR	Short Circuit Ratio

SIR	Synchronous Inertial Response
SNSP	System Non-Synchronous Penetration
SONI	System Operator for Northern Ireland
SPS	Special Protection Scheme
SRP	Salt River Project
SRS	Supplemental Reserve Service
SVC	Static VAR Compensator
UFLS	under-frequency load shedding
VRB	Vanadium Redox Battery
VSC	Voltage Source Converter
VSPSH	Variable Speed Pumped Storage Hydropower
WAMS	Wide Area Measurement System
WECC	Western Electricity Coordinating Council
WTG	Wind Turbine Generator
WWSIS	Western Wind and Solar Integration Study

1 EXECUTIVE SUMMARY

1.1 Project Objectives

This report discusses Fast Frequency Response (FFR) as a potential mitigation option for the Australian Energy Market Operator (AEMO) in maintaining grid security during low inertia conditions on their system. Broadly, FFR is rapid injection of power or relief of loading that helps arrest the decline of system frequency during disturbances. (A more detailed discussion is provided in Section 2.1) The first part of this report focuses on the practicality, requirements and capabilities of various non-synchronous resources to provide FFR. In the second part, we examine system requirements and benefits of FFR, and provide discussion of considerations for procuring FFR services.

Primary Frequency Response (PFR) is used in this report to denote the relatively fast, autonomous reaction of system resources to change in frequency. In most power systems, the main contributor to PFR is the governor response of synchronous generation. In the National Electricity Market (NEM), the 6-second Frequency Control Ancillary Service (FCAS) product is the means by which PFR is formally provided. However, throughout the illustrations and discussions, we have used a generalized representation of governor response[1].

To demonstrate the interplay of energy, speed, inertia, FFR and PFR, we provide illustrative simulations on an extremely simple model of the South Australian system under a variety of system conditions. One key element to maintaining system security in a power system with low inertia is to provide adequate arresting energy to stabilize frequency. This theme is carried throughout this work.

Broadly stated, AEMO's goal is to maintain a level of system security consistent with both the specific requirements of the Australian grid code (the National Electricity Rules), and with accepted industry practice. The primary focus of this work is aimed at technology options that have the potential to minimize or avoid involuntary under-frequency load shedding (UFLS), which starts as frequency decreases to 49 Hz. A further objective is to advance the overall security and robustness of the system, of which avoiding 'system black' events is a significant consideration. The scope of this report has the narrow constraint of only considering technologies *outside* of "conventional" solutions ("conventional" solutions might include addition of synchronous generation, new transmission lines, synchronous condensers or imposing operating constraints).

In the situation of South Australia (SA), there are credible losses, such as one circuit of the Heywood interconnector. Further, there are non-credible events, such as the loss of both circuits of the Heywood interconnector, which can result in the loss of up to (or sometimes exceeding) the maximum planned infeed of 650MW to South Australia, which needs to be

[1] This assumes a closed-loop response (i.e. controlled increase in power for a drop in frequency) that is representative of how equipment suppliers tune power plants. For example, we have used the IEEE standard model GGOV1 for the responsive generation in South Australia. The question of exactly how the present resources in the NEM, and particularly in South Australia, perform is of great interest, but is outside the scope of this work.

addressed separately. This is an extremely disruptive event, which has more severe impacts on the SA system as import levels increase. The event severity increases even further under low inertia conditions. While UFLS is allowed for these non-credible events, complete system black must be avoided.

The quantitative illustrations and many observations throughout this report are based on the situation in SA. SA was used as a case study for this analysis because the rate of change of frequency (RoCoF) has been identified as an immediate challenge in SA[3]. While created with the situation in SA as a significant consideration, the discussion is intended to provide more general insight, applicable to the entire NEM. Issues of economically maintaining good frequency response, with consideration of all available technologies to do so, is of interest for the entire country.

1.2 Key Findings

Several "unconventional" technologies can provide Fast Frequency Response

Section 3 of the report examines the capabilities of the following technologies to provide arresting energy for FFR:

- Wind turbines
- Lithium, Flow, and Lead-acid batteries
- Flywheel energy storage systems (inverter-interfaced and non-inverter systems)
- Supercapacitor energy storage systems
- Solar PV
- Load based resources
- HVDC transmission

All were found to have the potential to provide FFR, but their characteristics, opportunity costs, economics, and practicality varied. Some of these, e.g. batteries, flywheels and supercapacitors, would be new resource additions that would need to recover costs by providing FFR and possibly other services. Some of these, e.g. wind and PV, are being installed for other reasons but would need to be procured with FFR specifications which may have impacts on capital costs. Some of these, e.g., PV, may have an opportunity cost if they are pre-curtailed to provide FFR. Table 1 summarizes some key characteristics of the different resources.

Because FFR characteristics (speed, shape, duration) vary, this study created a new analytical method to compare the effectiveness of different types of FFR (see Section 7). With this new methodology, both planning and procurement can consider a range of technologies without being unduly prescriptive of performance characteristics. Discussion is provided throughout the report about physical limitations and practical considerations of specifying and obtaining different performance from these FFR technologies.

[3] https://www.aemo.com.au/-/media/Files/Electricity/NEM/Security_and_Reliability/Reports/FPSS---Progress-Report-August-2016.pdf

The application and opportunity cost for most of these technologies is straightforward. A notable exception is that inertia-based FFR from wind turbines (IBFFR) does not need to be pre-curtailed (or charged) to provide this service (although the energy recovery after IBFFR is deployed needs to be considered). Because the available arresting energy with IBFFR is quite limited, and because provision of this service by wind turbines has low operational and opportunity costs, it may be appropriate to handle this type of service separately from FFR technologies that can be sustained for longer periods. The aeromechanics of wind turbines makes the practical upper limit of IBFFR control for wind turbines about 10% of the power production. So, at an example level of 1000MW of wind power production, about 100MW of IBFFR should be achievable.

PV could also be an important FFR contributor. The possibility for utility-scale PV to provide low marginal cost FFR by taking advantage of emerging practice of relative short-term ratings of panels and inverters presents a potentially economic option which should be considered.

Both synchronous and asynchronous flywheels have beneficial attributes that should be considered. Synchronous flywheels (high inertia synchronous condensers) are a technically effective means of increasing inertia and support grid voltage and short circuit strength.

Many of these resources, including wind and PV, can also provide PFR (FCAS in the NEM). This can benefit system robustness. Provision of PFR (even for wind turbines, because PFR is not inertia-based) typically imposes a substantial operating cost on them (in the form of lost energy delivery, green certificates or other metrics associated with production), that needs to be addressed in the procurement process.

Low inertial conditions require fast responses

A critically important characteristic of these resources is response time. Lower levels of inertia on the system will result in higher RoCoF, i.e., for a given contingency, lower inertia means frequency will drop faster. In order to arrest the system before frequency reaches UFLS trip points, fast response time is required. Total response time is comprised of the following:

- Measure – Measure and identify frequency deviation and fast frequency decrease.
- Identify – Identify occurrence of severe event that requires FFR.
- Signal – Communicate action to be taken.
- Activate – Actuate the resource.
- Activate fully – Full response from resource.

A summary of response times for the technologies above is presented in Table 2.

Table 1. Summary of key metrics for each technology

	Wind Turbines	Lithium Batteries	Flow Batteries	Advanced lead-acid Batteries	Flywheel (inverter)	Flywheel (non-inverter)	Super-capacitor	Solar PV	Load Tripping	HVDC
Once triggered, time for full activation	0.5-5 s	Few – 40 ms	Few – 40 ms	Few – 40 ms	≤4 ms	As per synchronous devices	10-20 ms	100-200 ms	Depends on load	50-500 ms
Short-run marginal cost for FFR	Small for inertia-based FFR; opportunity cost with curtailment	Very small (standby losses)	Very small (standby losses)	Very small (standby losses)	Very small (standby losses)	Small (standby losses)	Very small (standby losses)	None, if overload. Opportunity cost if curtailed.	Depends on load	None; Opportunity cost if curtailed.
Maturity	Commercial, but limited deployments: First generation offerings of 'synthetic inertia'	Rapidly moving to mature	Relatively immature	Based on mature technologies	Commercial, but limited deployments	Based on relatively mature technologies; rapidly evolving designs	Relatively immature; evolving rapidly	Relatively immature for controlled active power applications	Relatively mature; technologies for widespread tripping of many small loads emerging	Mature, for triggered response functions
Other important considerations	IBFFR is energy limited	Capital costs dropping rapidly	Energy rating highly scalable	Can be attractive for applications with low cycling (i.e. contingencies)		Contributes short circuit strength. Transient stability behavior different.	High power density; tend to be specialized, small applications now	May depend on relative rating of panels and inverters		.

Table 2. Summary of response times for various non-synchronous resources

	Measure	Identify	Signal	Activate	Activate Fully
Options for external detection and signaling					
Direct detection	≤2-3 cycles 40-60ms		~1 cycle 20ms		
RoCoF detection w/PMU	~2-3 cycles 40-60ms		~1 cycle 20ms		
Local RoCoF/Frequency measurement	≥5 cycles 100ms		nil		
FFR options that require external detection and signaling					
Wind turbine with inertia-based FFR				~2 cycles 40ms	~500ms
Lithium batteries				10-20 ms	
Flow batteries				10-20 ms	
Lead-acid batteries				40 ms	
Flywheels (inverter)				≤4 ms	
Super capacitor				10-20 ms	
Solar PV				100-200 ms	
HVDC				50-500 ms	
FFR options that detect, signal and actuate					
Flywheels (non-inverter)	instantaneous				N/A
FFR options that signal and actuate					
Load			2-3 cycles 40-60ms	Depends on load	

The main challenge is to *quickly and accurately* identify a severe event that requires FFR. Complicating this step is the fact that directly after an event, frequency varies spatially. So, while one part of the grid may perceive a severe event, another part of the grid may not. Additionally, triggering too much FFR may have adverse consequences, as this work demonstrates. Table 3 summarizes the tradeoffs between different types of detection and topology of FFR response.

Table 3. Tradeoffs in topology and detection mechanism

	Triggered by local RoCoF measurement	Triggered by direct event detection
Few, large resources	Pros: Cheaper topology Cons: High risk of false triggering	Pros: Low risk of false triggering Cons: Only applies for specific events; Moderate cost for communications
Many, small resources	Pros: Reduces consequences of false triggering Cons: Expensive detection	Pros: Low risk of false triggering Cons: Expensive communication

Another difficulty is that very fast measurements may misinterpret transients, switching operations or other actions that are *not* severe events as reason to trigger. Risk of false triggering is mitigated by longer periods for measurement and identification, but this comes at the expense of FFR activation time. So, FFR response time is critical, however a balance between making high fidelity decisions to act and speed is needed Fortunately, it turns out that FFR needs to be fast but not incredibly fast. Simulations in Section 7.2 show the efficacy of injections of FFR after a contingency event with varying delays in response time. FFR needs to be started well before UFLS or the occurrence of the frequency nadir[4]. Analysis presented suggests that total response times on the order of one quarter to one half second are sufficiently fast. It is non-intuitive, but extremely fast FFR is less effective. If it is too fast, then it interferes with and stifles full PFR response. Part of the planning process can include fine-tuning the response time of FFR, thereby improving the efficacy of the FFR for critical conditions.

Fast Frequency Response, Primary Frequency Response, and Inertia all interact

There is a delicate interplay between FFR, PFR and inertia. The primary function of FFR is to arrest the frequency decline and "buy time" for PFR[1] to act. The amount of FFR needed and its efficacy is closely tied to the amount and quality of PFR available. For example, faster PFR will reduce the amount of FFR required at any given level of inertia; however at very low levels of inertia, conventional PFR (from synchronous generation) has limited ability to provide arresting

[4] Nadir is the minimum frequency during the swing from a system disturbance (see Section 2.1).

energy fast enough. Similarly, withdrawal of FFR should avoid abrupt steps, and should be coordinated with the PFR.

Determining system needs and practice for FFR requires detailed simulations. This report uses power system modeling to examine the interplay of inertia, FFR and PFR in detail. As noted, for this work, a dynamic model of the South Australia power system was developed, with input from AEMO. The model was designed to be relatively simple, but tuned to give the same response as measured in South Australia for one event[5]. The model was exercised to investigate the impact of speed and energy injected on FFR efficacy.

The largest factor in determining response is the size of the event, i.e., the amount of power infeed lost in the disturbance. Figure 1 shows the FFR requirement as a function of contingency size for two different inertia levels. The efficacy of adding inertia versus adding FFR is given by the difference between the red and blue curves. For example, at 450MW import (as indicated by the blue arrow), it takes about 120MW more FFR to avoid UFLS when the system inertia is 1000 MW-sec than it does when the system inertia is 9000 MW-sec. Under these conditions, the FFR must act within about 250ms for the 1000 MW-sec condition, whereas this result for the 9000 MW-sec is based on 500ms action.

The economics can also be estimated from this. The capital cost to meet the marginal requirement for FFR is the difference in cost of (about) 330 MW of 500ms response FFR compared to the cost of 450 MW of 250ms response FFR. A very rough estimate for 120MW dedicated energy storage device(s) would be A$120M (This is a high estimate, since at least some of the necessary FFR could come from relatively less expensive resources, like IBFFR from wind generation). The cost to *add* inertia (in the form of synchronous flywheels) is on the order of A$35/kw-sec, so 8000 MW-sec, would be roughly A$280M. Note that this exercise assumes a degree of linearity that may be at odds with the exact details of the system performance. While this is a simplistic model with rough estimates, it suggests that new infrastructure to provide inertia may need additional value streams to recover costs.

Getting the right amount of FFR becomes more important as inertia decreases. The amount of FFR that is needed will be dominated by (1) the amount of load shedding that is allowed for any particular operating condition and (2) by the size of the event. The ability to correctly shed the right amount of load based on local measurement of frequency and RoCoF declines with inertia. Presently observed levels of inertia in South Australia make autonomous UFLS of questionable robustness. In particular, the use of highly distributed loads to provide FFR is likely to be subject to the constraints and uncertainties of rapidly measuring RoCoF and frequency.

---

5 System separation on November 1, 2015.

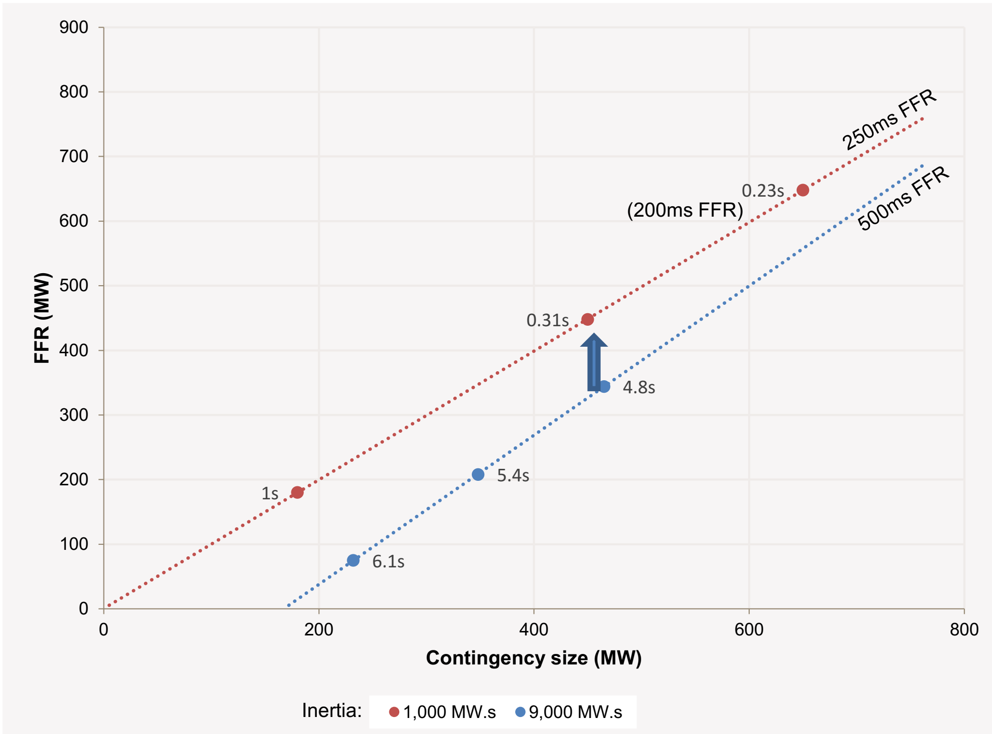

Figure 1 FFR Requirement vs. Contingency Size (for different system inertia). Data labels show the time to the nadir in seconds.

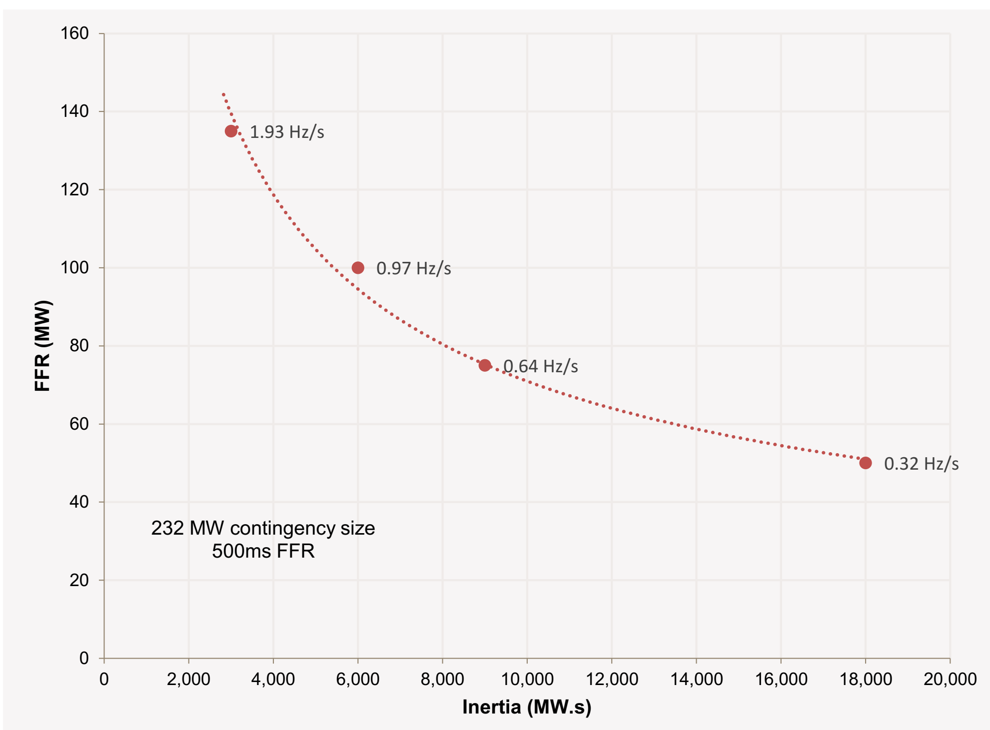

Figure 2 FFR requirement vs. Inertia for given disturbance size (with initial RoCoF noted)

Simulations show how much FFR is required to avoid UFLS (49 Hz) for various levels of inertia in South Australia (see Figure 2). For this case, the AC import that is lost when the Heywood interconnector is opened is 232MW (approximately the amount lost in the November 1, 2015 event). As the system inertia decreases from 9000 MW-s to 3000 MW-s, the amount of FFR required increases by about 60MW.

Costs of FFR should be compared to the cost of alternative operational strategies

In order to determine the relative marginal economic efficacy of adding FFR compared to imposing operating constraints in the form of constraints on South Australia inertia and imports from Victoria, detailed comparative production simulations are required. Unfortunately, imposing dynamic constraints on production simulations can be challenging. The results of dynamic simulations need to be mapped to static boundary conditions that can be modeled in a production simulation.

There is industry precedent for this practice. For example, (in North America) transient stability derived path ratings are often mapped into (so-called) nomograms. In Figure 3, a set of nomograms show how a production simulation could impose the interrelated constraint of import limit on the Heywood interconnector with committed synchronous inertia. The system is assumed to be constrained by "available response". We introduce this concept here, because from a variable cost perspective (i.e. the constraints of a production simulation) willingness to shed load by UFLS (or other means) and the amount of FFR available are equivalent and additive.

The different traces are measures the sum of the allowed load shedding plus the amount of FFR provided. For clarity, consider the blue "no response" trace. If we suppose that no load shedding is allowed and that no resources are available to provide FFR, then this line represents the loading constraint on the AC interface. At minimum inertia (of 1000 MW-s), zero power can be imported. As inertia increases (e.g. by commitment of synchronous generation in South Australia), then import limit eases, up to about 165MW at 9000MW-s. Now consider the red "200 MW" line. If 200 MW of FFR is available, then at minimum inertia, the AC import limit is 200MW. At 9000MW-s, the import limit is 330MW. But, 200 MW of load shedding would produce the same result, as would (say) 100MW of FFR and 100MW of load shedding. This is a production costing constraint. Other measures, including those that might require capital expenditures (e.g. reactive compensation, or improved controls), could be needed, but they have no bearing on the variable cost of operation. Using pairs of simulations with this type of constraint allows a meaningful calculation or comparison of what FFR (or UFLS) is "worth", from a variable cost of operations (and emissions) perspective.

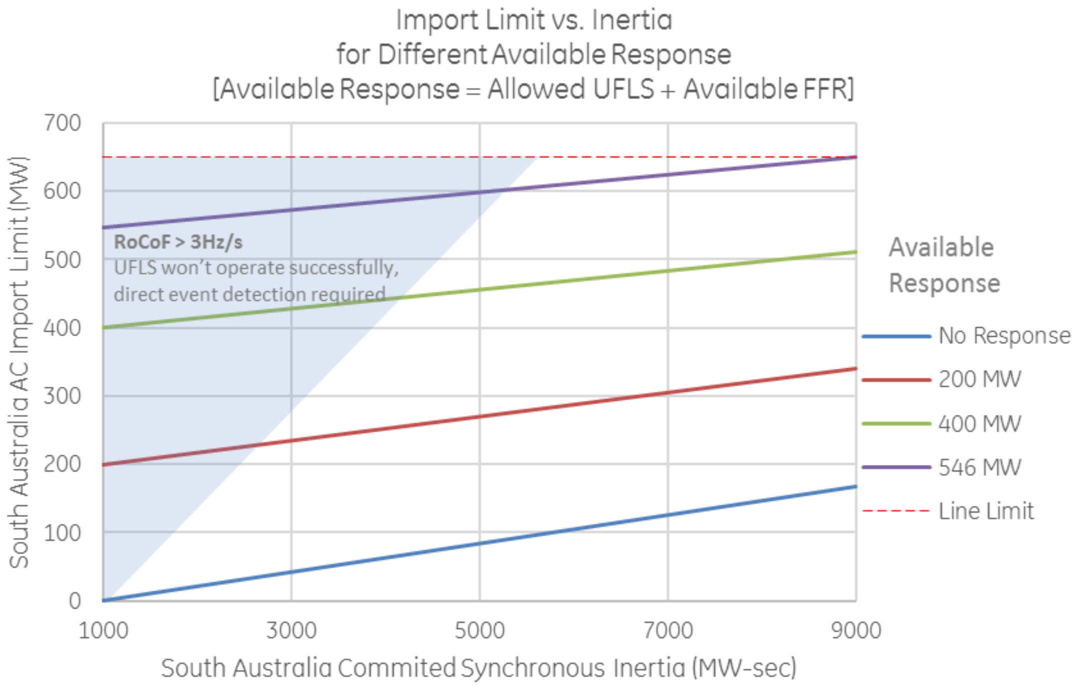

Figure 3 FFR Requirement Nomogram for Production Simulations of South Australia

A properly designed production simulation will, at each operating interval (e.g. hourly), select the most economic combination of committed generation and import (i.e. committing more synchronous generation to adds inertia, raising the import limit as much as 175MW). If the threshold is raised to (say) 200 MW of load shedding allowed or 200 MW of FFR available, the red trace will apply instead. A production simulation with this eased constraint will result in lower overall variable operating cost (VOC) of energy (on an annual basis, not necessarily at each hour); it may result in lower carbon emissions as well. The difference provides a calibration on the economic efficacy of relieving the constraint. That is, the results will show what could be spent (annualized capital cost) on FFR, or what would be saved by allowing more load shedding. This approach would allow for the evaluation of added inertia as well; i.e. if dedicated flywheels were added to increase inertia all the time, the constraint would shift to the left, rather than up. A similar comparison of annual VOC would show the energy and emissions value.

The primary method to manage South Australia islanding should be by Special Protection Scheme

This report finds that the most effective solution for managing the loss of the Heywood interconnector will be a load tripping Special Protection Scheme (SPS), that is continuously updated (e.g. by SCADA), armed and designed to interrupt load and trigger other discrete FFR resources. The scheme should include some highly responsive, closed-loop continuously acting FFR in South Australia. These resources need to be available (i.e. on line and not in limits[6]) during periods of risk. The amount armed by such an SPS should be an amount close to equal to the size of the event. The amount required will decrease somewhat with increased inertia (see nomogram in Figure 3). Considerable further technical and economic (and societal) analysis is required to establish the right mix of new FFR technologies and direct load control for the SPS.

It is worth noting that UFLS does not help prevent separation due to overload of the interconnector (e.g. the September 28, 2016 event). A suitably designed SPS would help.

Maintaining stability is more than just matching load and generation

It is critical to ensure system stability is maintained, especially in South Australia, while examining options to manage under-frequency. That is, the South Australia system must be able to transition to an island without loss of synchronism within the South Australia system. The reactive power requirements of the system must be met. Good reactive power supply and control can play a critical role in preventing an island from forming in the first place. The system must have positive damping. Present technology inverter-based resources like wind and solar PV tend to have better behavior with respect to power system oscillations. But system dynamics change, sometimes dramatically with altered flow patterns and commitment of synchronous generation. Inverter based devices must have sufficient short circuit strength to maintain control stability. These considerations are outside the scope of this report.

Further, the stability issues associated with islanding ought not to be addressed separately. For example, as recent events have shown, large upsets, such as loss of generation within South Australia, must also have stable outcomes. The best solutions to the islanding challenge should be designed to produce benefits for other stability risks (beyond islanding) as well.

This is new territory with respect to system dynamics. The scale of the events in South Australia, i.e. the amount of power that can be lost when the interconnector is heavily loaded, relative to the total load being served in the receiving system, is extreme relative to most bulk power systems. This report attempts to distill a large body of state-of-the-art technology developments but further evaluation, testing, and analysis is required to ensure system security in low inertia systems. Details matter, and the model used in this study is not a substitute for real simulations with high fidelity models.

Zero synchronous generation is a near future condition that needs to be considered. Technology, understanding and operational practice for zero synchronous operation are outside the narrow scope of this FFR investigation. South Australia should, *for now*, proceed with the understanding that there should be zero expectation today that it can survive, islanded from the rest of the NEM, without at least some synchronous equipment in operation.

---

[6] To get the desired response from a particular resource (say a battery energy storage system for example), it must be energized and synchronized with the grid ("on line") and it must have the capability to respond ("not in limits"). So, to continue the example, if the BESS were running flat out, say, selling power into the grid for arbitrage revenue, it would not have capability to provide FFR too at that instant of time. This simple, but important, constraint can be easily overlooked.

While the state-of-the-art is rapidly evolving, this condition will require significant work in terms of infrastructure, understanding and adaptation of operational practice.

Load must continue to be a resource

Load tripping or control will continue to be part of risk management for South Australia. For this discussion, fast voluntary disconnection of load, i.e. load tripping, should be considered a legitimate type of FFR. This type of load control is a technically (and often highly cost) effective option. This is a hugely important distinction from the type of autonomous, localized and involuntary load shedding (i.e. UFLS) that is to be minimized. The presumption, from a procurement perspective, is that load participation in FFR would be voluntary and compensated. However, from a security perspective, the procurement aspect is a separate discussion from that of effectiveness, economy and robustness.

For separation (islanding) of the South Australia system from the rest of the NEM – that is trip of the Heywood interconnector or similar events (e.g. high loss of generation within South Australia), economic solutions to maintaining acceptable performance will likely require that service to some loads in South Australia be interrupted.

Methods of load tripping and control will need to evolve from present practice. Two important factors will drive the evolution of load tripping from present day practice of actuating utility (e.g. distribution feeders) towards more sophisticated methods. First, with the continued growth of highly distributed generation, mainly PV, opening feeders disconnects generation as well as load. It is entirely possible that the system, especially South Australia, will see periods of zero net load in future. Second, as observed above, the requirement to have ever more precise, real-time information about the exact amount of load reduction that will result from each action grows with dropping system inertia.

Conventional UFLS cannot be relied upon to manage extremely high RoCoF. Historically, South Australia has depended on UFLS that acts at or below 49Hz as a critical line of defense to "save" the system. "Save" in this sense means that the frequency swing should not go below 47Hz, and the system should not go black. For the low range of inertia considered in this work (i.e. down to 1000MW-s), this almost certainly optimistic. UFLS probably cannot be relied upon to have the desired outcome for the very high RoCoF levels that will accompany simultaneous low inertia and large disturbances (e.g. high loading on the AC interconnector). Therefore, resources addressing low inertia conditions should be mostly, if not all, triggered by direct event detection, rather than local frequency or RoCoF detection. Nevertheless, UFLS can still continue to provide an important contribution to system security. It is envisioned that conventional UFLS will continue to be used, and will serve to provide important protective functions.

To design a robust UFLS solution, a detailed examination of the tradeoffs discussed here is needed. Further work needs to go into:

- Understanding the impacts of false triggering;
- Understanding how easily severe frequency events in South Australia can be discerned from other disturbances in South Australia; and

- Understanding how local frequency varies across South Australia due to a severe event.

FFR procurement should consider both capital and variable costs

It is first important to note that physics should drive FFR procurement: maintaining the security of the NEM, especially the South Australian system, is becoming increasingly challenging. This is emphatically a power system dynamics problem that requires good utility engineering practice. Market constructs must be founded on this, and must have sufficient flexibility to adapt to what will inevitably be a rapidly changing situation – both in terms of available technology and details of the grid dynamics.

Provision of FFR from most resources that might be built for the specific purpose of providing that service involves significant investment in physical equipment. For dedicated resources like batteries, flywheels, etc., the capital investment is the dominant cost. The variable cost of providing FFR is usually relatively low. Other resources, which might be able to provide FFR *in addition* to their main function (e.g. IBFFR from Wind, or FFR from temporarily over-driven solar PV), may be able to offer the service after a relatively modest capital investment. Some resources, like batteries, might also incur opportunity costs if they are performing other market functions, such as energy arbitrage[7]. Design of a market for procurement of FFR needs to consider these realities.

Procurement by an annual (or periodic) clearing market (i.e. winning resources are obligated to provide FFR service when called upon, in return for a fixed revenue stream, plus energy net) can be effective. Tendered procurement, in which a call is issued for competitive supply of FFR, can also work. Variable (e.g. Day-Ahead clearing) markets for FFR may face challenges for price formation and stability.

1.3 Summary of Recommendations

- A comprehensive effort to specify and design a Special Protection Scheme to address separation of South Australia from the rest of the NEM should be a high priority.
- Properly structured production simulations, that can capture these new constraints and consider provision of FFR and FCAS from a full range of technologies should be run. Economic evaluations, considering both capital and operating costs, are needed to provide meaningful comparisons between options. FFR options should include the technologies outlined in this report and conventional resources, as well as altered, more conservative operating strategies and relaxed performance standards.
- Model validation for evaluation of the dynamic performance of the NEM is needed, with attention to whether the available models provide adequate fidelity for simulations of the more extreme conditions under consideration. Performance of all important elements, including renewable generation and all resources providing FCAS must be considered.

[7] Energy storage may participate in energy markets, buying at low prices and selling at high prices. Reserving power rating for FFR could result in lost revenues.

- Detailed dynamic analysis of the NEM is needed. Illustrative cases and relationships developed with the very simple model of the South Australia system in this report provide guidance for relationships to be more formally calculated.
- Consider system solutions that will maintain a minimum level of inertia and that will contribute other performance benefits as well (i.e. improved transient, frequency and voltage stability).
- The possibility for utility scale solar PV to provide low marginal cost FFR by taking advantage of emerging practice of relative short-term ratings of panels and inverters should be investigated (this is an industry-wide issue and recommendation).

2 INTRODUCTION TO FFR CAPABILITY AND TECHNOLOGY

Power systems are designed to maintain security of the system in the event of a contingency, such as the loss of a large generator. This loss is invisible to the end-users, who see no interruption in service or quality. In general, each interconnected power system should have enough arresting power so that UFLS does not occur for most credible events[8]. UFLS, in which blocks of firm load[9] are disconnected from the grid to keep the rest of the grid operational, starts at 49 Hz in the NEM.

Different power systems use different terminologies and different services with various characteristics to arrest and then stabilize and recover frequency. The basic concepts remain the same.

2.1 Fundamentals of Frequency Response

Key elements of frequency response as presently viewed by most of the power industry, are shown in Figure 4. When generation is lost, frequency across the interconnected power system will decline. Inertia from the rotating mass of generators and induction motors that are online will determine how fast frequency falls immediately following the event. (This is called the initial RoCoF). The first few seconds, up to the time of a frequency minimum (the frequency 'nadir') is called the arresting period, and is shown in blue in Figure 4. The inertia of the system is not a control action, but rather a physical characteristic of the rotating mass of whichever synchronously connected devices are online at the time of the disturbance. Without further action, the imbalance between load and generation will result in continuing decline of frequency. Thus, inertia alone is insufficient. Historically, arresting the decline of frequency results from the deliberate control actions taken autonomously by elements of the power system: that is mainly individual generators with active speed governors and room to increase output (headroom). This PFR will normally stabilize frequency in the seconds to tens of seconds time-frame. While the figure shows this response starting immediately at the time of the disturbance, it is important to note that it normally takes a second or so before any additional power is injected to the grid due to governor action. And that it normally takes several seconds, up to tens of seconds, before typical turbine-generators fully response to the frequency error. In the NEM, there are a range of frequency control ancillary services presently procured, including a 6 second product and a 60 second product. For the purpose of this discussion, we are primarily concerned about the behavior of the system from the initiating event up to about 20 or 30 seconds after the event. Therefore, it is (our understanding that) the 6 second product that provides most of the PFR. As the size of system disturbances increases, and as the amount of inertia decreases, the amount and speed of PFR needed to

[8] Interconnected power systems are not designed to withstand loss of every generator, but rather are designed to withstand loss of large, credible contingencies. For the NEM, loss of one circuit (e.g. 1 of the Heywood interconnector circuits) is considered credible. *Simultaneous* loss of 2 circuits (e.g. both circuits of the Heywood interconnector) is not considered credible by NEM rules. Similar rules exist elsewhere; for example, in the US Western Interconnection, this "design-basis event" is the loss of 2 large nuclear units.

[9] It is noted that with the significant increase in behind-the meter distributed energy resources, "firm" load may now include generation. This new twist on UFLS schemes is important to consider for those regions where distributed generation may be significant, and deliberate disconnection of that generation may have systemic effects counter to the intent.

arrest the system frequency decline before reaching a frequency that results in involuntary disconnection of customers, i.e. 'involuntary UFLS', increases. Following the frequency nadir is the rebound period, shown in yellow in Figure 4. The PFR will allow the frequency to settle at a new point based on the load/generation balance.

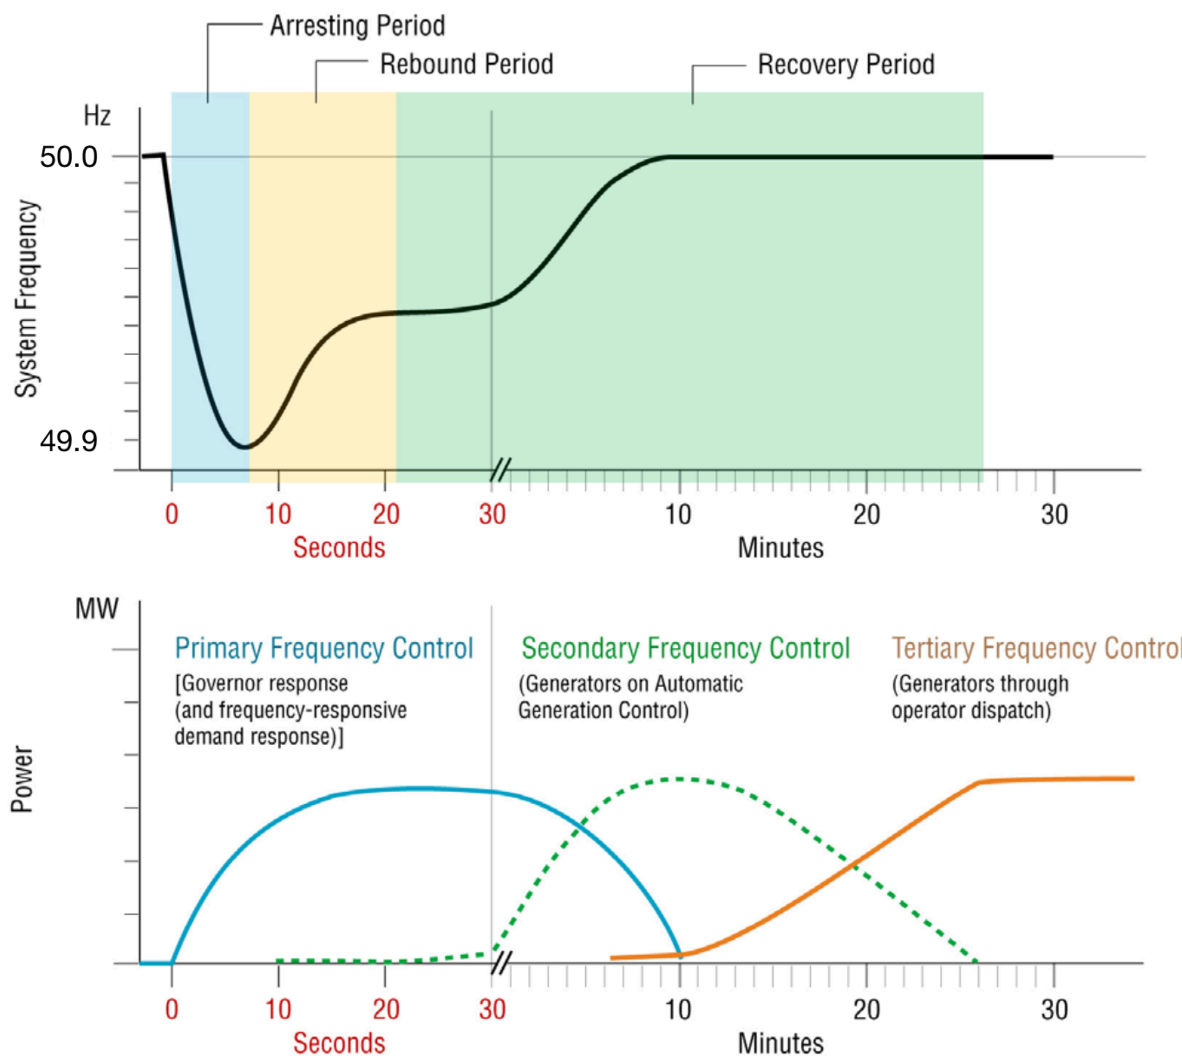

Figure 4 Frequency response to a loss of generation event[10].

The growing need for speed of response has given rise to consideration of a new, distinct service: FFR. FFR is similar to PFR but acts much faster, providing power during the arresting phase, with the specific objective of providing arresting power before the frequency nadir. PFR

---

[10] Adapted for 50 Hz system from J. Eto, et al. "Use of Frequency Response Metrics to Assess the Planning and Operating Requirements for Reliable Integration of Variable Renewable Generation", LBNL-4142E, Dec. 2010. http://www.ferc.gov/industries/electric/indus-act/reliability/frequencyresponsemetrics-report.pdf

and FFR both help to arrest frequency and interact with inertia to determine the frequency nadir. FFR will also contribute to establishing the settling frequency [11], if the FFR is sustained past the time of the nadir into the rebound period. In the most common definition, both PFR and FFR are autonomous controls that act based on local conditions, that is, they respond to quantities like local frequency or machine speed) that can be measured at, or very close to, the equipment providing the service. An example of the interplay between PFR and FFR in ERCOT's proposed ancillary services is described in the text box.

Secondary frequency response will increase output during the recovery period, shown in green in Figure 4 to return frequency to nominal. Secondary frequency response occurs in a few minutes. Unlike PFR and FFR, secondary response is often a centralized control action, with the central energy management system of the grid operator sending power instructions to individual resources (mostly participating generators) in response to deviation of system frequency and interconnector flows. Finally, tertiary reserves will be dispatched up to allow frequency regulation units to be dispatched back down and be available for the next disturbance. Tertiary reserves often have market objectives, i.e. are "economic redispatch".

Frequency response services are also used to mitigate over-frequencies, which could occur if a block of load is dropped. In this case, generators respond by decreasing output. This is discussed further in Section 4.4. While this is an important response and many of the technologies in this report can respond well to both under- and over-frequency events, the bulk of this report is focused on under-frequency mitigation which is a harder problem to solve and the greater concern of AEMO.

[11] "Settling frequency" is the period during which the PFR has stabilized the frequency from an event, but before the secondary response, i.e. REG (or 60 second FCAS) begins to restore frequency to nominal 50Hz. It is represented by the flat section between 20 and 60 seconds in the figure.

Future Ancillary Services in ERCOT

ERCOT is in the process of redefining their ancillary services (AS) as the ERCOT system evolves[12] (changing generation mix, some of the new resources have new challenges and some of them have new capabilities for AS, new regulatory requirements). Their proposed new services include Synchronous Inertial Response (SIR), Fast Frequency Response (FFR), Primary Frequency Response (PFR), Up and Down Regulating Reserve (Reg), Contingency Reserve Service (CRS), and Supplemental Reserve Service (SRS) as shown below.

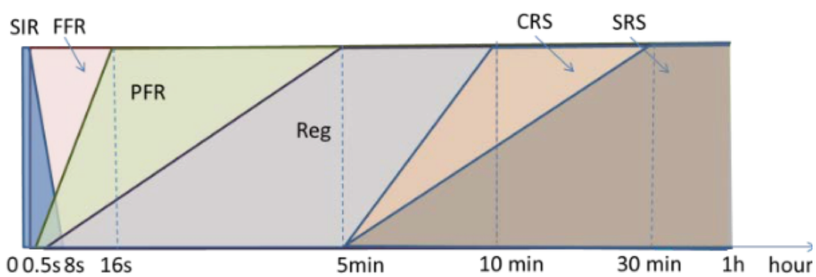

SIR is an instantaneous response that acts to decrease RoCoF, enabling time for PFR to arrest frequency. FFR is a fast-acting response: the time from frequency meeting a threshold to full response from the resource is 0.5 seconds. This is sustained for 10 minutes or until ERCOT recalls the resource, whichever is less. FFR must then be restored in 90 minutes and ready to be deployed again. FFR would be a new service, although 1400 MW of current Responsive Reserves from load resources meet the FFR requirements. PFR is used to arrest frequency decay and provides continuous full proportional response to frequency deviations. ERCOT has studied the interaction between FFR and PFR. An example of how FFR may reduce PFR requirements is shown below.

System Load	FFR	PFR(5% Droop)
<=35000 MW	840 MW(30% 0f 2800)	1960
<=35000 MW	0 MW	4480 MW

Security constrained economic dispatch automatically dispatches generation every 5 minutes to balance the system. Within those dispatch intervals, Reg follows an AGC signal every 4 seconds to provide fast balancing of the system. CRS ensures that within 15 minutes of a disturbance, frequency can be restored. CRS needs to provide 95% of full response within 10 minutes after receiving ERCOT signal. Within 90 min of frequency restoration, PFR and Reg are restored and ready to deploy again. Finally, SRS is a longer term product used to compensate for forecast errors.

[12] The ERCOT process is on-going. Changes have not yet been approved. http://www.ercot.com/committee/qmwg

2.2 Basic Components of Fast Frequency Response

Initial RoCoF is the rate-of-change-of-frequency immediately after the system is unbalanced by a disturbance like loss of generation or islanding. Initial RoCoF is related to synchronous system inertia and the size of the contingency according to the following equation (assuming load dampening to be zero):

$$RoCoF = \frac{df}{dt} = \frac{\Delta P}{2H} \times f_0$$

where

$$\Delta P = Size\ of\ contingency\ (MW\ lost)$$

$$H = System\ inertia\ \ (MW - seconds)$$

$$f_0 = Frequency\ at\ the\ time\ of\ disturbance\ (Hz)$$

$$\frac{df}{dt} = Rate\ of\ Change\ of\ Frequency\ (Hz/\sec)$$

In the instant following the disturbance, no control actions take part. Notice that the equation does not include any terms other than the size of the event and the inertia. Therefore, for systems where high RoCoF can drive adverse behavior, such as tripping of generators and loads, high system inertia is desirable. There are two ways to manage frequency response:

1. **Slow RoCoF down**: Synchronous generators provide inertia that decreases RoCoF. Governors on conventional generators act fairly slowly (on the time frame of seconds to minutes) to increase output to arrest and recover frequency. A conventional enhancement would be to increase grid inertia using synchronous machines (e.g. synchronous condensers) to slow RoCoF to allow time for other responses to activate. Synchronous machines do this autonomously and with no need for 'detection' or 'communication'. This is how the power system is managed at present and therefore AEMO has directed that this not be the focus of this report.

2. **Respond more quickly**: Another option is to mitigate the effects of fast initial RoCoF by rapidly injecting arresting energy into the grid. Note that arresting energy could be provided by increasing generation or decreasing load. Both provide the same function of bringing system generation and load into balance. The objective of this study is to examine what kind of very fast frequency response can be achieved using advanced features from new inverter-based technologies such as (but not exclusive to) wind, solar, batteries, flywheels. This option requires measuring frequency and identifying the problem, communicating to the resource that a response is needed, and activating the resource. This approach is the focus of this report.

2.2.1 Discussion of South Australia Situation

The South Australia situation is somewhat unique and is significantly more severe than what has been experienced in other regions of the world. South Australia demand ranges from 800 to 3400 MW. This is demand on the grid, and excludes the approximately 680 MW of distributed behind-the-meter PV. The Heywood AC interconnector has a maximum transfer limit of 650 MW for flows in or out of South Australia, and is the region's only synchronous interconnection with the rest of the NEM. At maximum demand, loss of Heywood means a 19% loss-of-infeed event (650 MW/3400 MW). At minimum demand, loss of 650 MW constitutes an 81% loss event (650 MW/800 MW).

Table 4. Time window for arresting energy before UFLS triggered

RoCoF	Time before UFLS triggered	Cycles[13] before UFLS triggered
1 Hz/sec	1000 ms	50
2 Hz/sec	500 ms	25
3 Hz/sec	333 ms	16.7
4 Hz/sec	250 ms	12.5

Two different types of events have distinctly different performance objectives[14]:

(1) If a single generating unit trips, frequency must be maintained within the containment band (49.5Hz) and the South Australia system should not experience involuntary load shedding.

(2) Simultaneous trip of both Heywood circuits should not cause full system black. Load shedding is allowed. Trip of one circuit, when they both are in service before the event, cannot result in tripping of the other circuit.

At high levels of loading on the Heywood interconnector, the simultaneous loss of both interconnector circuits is on the extreme edge of typical utility practice. This situation is quite challenging even before inverter-based renewables are added to the system. The relative size of these loss-of-generation events is more closely akin to that faced by large industrial facilities with internal dedicated generation (microgrids).

Calculations by AEMO have shown that if the Heywood interconnector is loaded at 650 MW and both lines trip, the RoCoF can exceed 4Hz/s. Because UFLS kicks in at 49Hz, this gives the system less than 250 ms (see Table 2) to arrest frequency before UFLS starts. The ability of conventional UFLS relays to successfully disconnect the right amount of load, fast enough to avoid unacceptable results (i.e. frequency below 47Hz and/or system black) is, at best, not

13 In 50 Hz system.

14 We use "objectives" here. It is outside of the scope of this work to either interpret the regulations or establish precise obligations or ability and jurisdiction to act to satisfy objectives.

assured. AEMO is therefore examining what options can arrest frequency in this time frame to ensure the RoCoF is slowed sufficiently for the UFLS to operate successfully.

This is pushing the envelope of existing requirements in other jurisdictions. ERCOT's FFR requires 500 ms; EirGrid/SONI's FFR proposed requirement is 2 seconds; UK National Grid's Enhanced Frequency Response requirement is 1 sec[15]. (See box in Section 3.3.) The time for which these various FFR requirements must be sustained varies, but the intent of all of them is that the response last at least until the frequency nadir. Some pilot projects are demonstrating fast response times with various technologies and these will be discussed in the subsequent sections.

The overall system solution may include a combination of several of these FFR inverter-based technology solutions together with some conventional synchronous solutions, and may also include accepting some additional risk. But because of the severity of the challenge, and because AEMO foresees a significant amount of new inverter-based generation being installed in the near future, low cost solutions may include provision of these types of system services from some of these new inverter installations.

It is important to view this issue with a realistic perspective. The worst-case events for the South Australia grid, i.e. those that occur at high loading on the Heywood interconnector, are significantly more severe (with much faster RoCoF) than what other similar grids are facing. The solutions to the events are beyond existing experience, and may be beyond the existing state-of-the-art applications of new technologies. New and innovative approaches will be required. For this discussion, fast disconnection of load, i.e. load tripping, should be considered a legitimate type of FFR. The concepts described in this report will each have some capability to contribute arresting energy that improves the frequency performance of the grid. The ultimate challenge is to balance the quantity and cost of the FFR resources with the achievable improvement in grid frequency response. The baseline against which such solutions are measured is the cost of reducing imports on the Heywood interconnector and running enough frequency responsive generation (which must be available) to avoid unacceptable performance. Today, frequency responsive generation means existing synchronous generation, but should include the possibility of PFR from wind and solar generation as well.

Later in this report, some additional context for economic evaluation of operating costs and benefits is provided, including participation by wind and solar generation. Involuntary load shedding for extreme but rare events may continue to be part of a practical solution, provided there are provisions for rapid restoration of load that was disconnected.

2.2.2 Response Trigger Options

There are two ways to detect the need to deploy FFR:

1. **Direct event detection** – Detect the specific condition of the disruption (e.g. the relay action that results in losing the AC link in South Australia) and having a direct transfer trip scheme inform the resource(s). This can be done quickly (on the order of a few

[15] National Grid, *Enhanced Frequency Response, Frequently Asked Questions Version 5.0*, 29 Mar 2016.

cycles) but it requires dedicated, fast communications and it only addresses the specific contingencies within its design criteria (i.e., if something else causes a frequency event, the FFR won't trigger).

2. **Frequency and RoCoF detection** – Detect the frequency deviation and high RoCoF. There are promising new technologies that claim to be able to do this very quickly. Accuracy and, especially, false triggering may still be an issue when attempting to measure frequency and RoCoF very quickly after a major system fault, however, as discussed below. Further, this approach has the limitation of only being applicable to frequency events (and not, for example, excess interconnector loading that might cause an island to form or other problems to evolve).

There are common elements to determining the potential for very fast response times. Figure 5 shows that the response time is the total of these actions:

1. **Measure** ($T_{Measure}$) – The controller must first measure and identify the frequency deviation and high RoCoF. It is important to note that *local* frequency may differ from *system* frequency directly after a disturbance and implications of this are discussed below.
2. **Identify** ($T_{Identify}$) – Identify the occurrence of an event that requires FFR. This means that the scheme must distinguish between a frequency event (for which action is needed) and other events, such as a fault that successfully clears, for which FFR is not needed.
3. **Signal** (T_{Signal}) – The controller must then communicate what action should be taken. Communications latencies are too slow to allow for most multi-unit plant-level controls to provide FFR. Instead, FFR controls may need to be implemented at the individual unit-level in order to respond quickly enough. Algorithms to control the response will also need to be fine-tuned to be very fast. This makes the use of distributed energy resources (such as rooftop PV or small loads) difficult because detection and controls will need to be distributed (in each individual PV inverter) and that cost may be prohibitive. However, there may be 'nearly instantaneous' methods of communications, such as powerline carrier approaches, that may reduce signal latency times to workable levels.
4. **Activate** ($T_{Activate}$) – The controller then actuates the technology. This portion of FFR performance is the time it takes for the technology to start acting (i.e., initiate its response).
5. **Activate fully** ($T_{ActivateFully}$) – The final measurement of FFR performance is not solely related to time: The objective of FFR is to *arrest* the drop in frequency. This is accomplished by delivering *energy* to the system. Thus, an appropriate measure of performance is based on how much energy is delivered in the required time. Here we are departing from some other proposals that focus on the rate of rise of the power injection. We note the three different subregions of response in the figure, with different hatching, and provide additional context for this approach in the subsequent parts of the report. Response times of the technologies vary as described in Section 3. A method for looking at the required response time and energy of dissimilar resources is presented in Section 7.

Total response time, $T_{TotalResponse}$, is then defined as:

$$T_{TotalResponse} = T_{Measure} + T_{Identify} + T_{Signal} + T_{Activate} + T_{ActivateFully}$$

Steps 1-3 are discussed below. Response times of various technologies (steps 4-5) are discussed in Section 3. There we present technological capabilities and performance. Performance testing has been documented for some projects, but it is noted that response times can vary. For example, depending on when in the control loop cycle the event occurs, the response could be signaled faster or slower.

In Figure 5, a simple illustration of the relationship between a frequency event and an FFR response is shown. At the beginning of the event (1 second), system frequency begins to drop at a RoCoF (as noted above) that is proportional to the size of the event, and inversely proportional to the system inertia. As the frequency drops the event must be detected, actions requested from FFR resources, they must respond by providing arresting energy to slow and stop the frequency decline. A primary objective is to ensure RoCoF is sufficiently reduced such that UFLS can operate successfully , so energy delivered before hitting UFLS (green shading) is most valuable. Even if UFLS occurs, FFR continues to arrest until the frequency nadir. This is the energy shaded in orange. After the frequency nadir, the FFR energy complements the restorative energy coming from the primary frequency response. This is shown in blue.

As the figure suggests, speed is important to avoid UFLS. However, we note that the power industry has not asked for response times that are as fast as what AEMO is discussing. Because of that, we note that technologies may not have been optimized for speeds less than several hundreds of milliseconds. However, there may be modifications or trade-offs that could reduce response times significantly. For example, an aggressive response from lithium batteries may be at the expense of increased thermal stress on the battery cells which could be mitigated by increased parasitic losses from increased cooling capacity or sacrificing lithium battery life expectancy. Rotating devices like flywheels and wind turbines will see higher torques, etc.

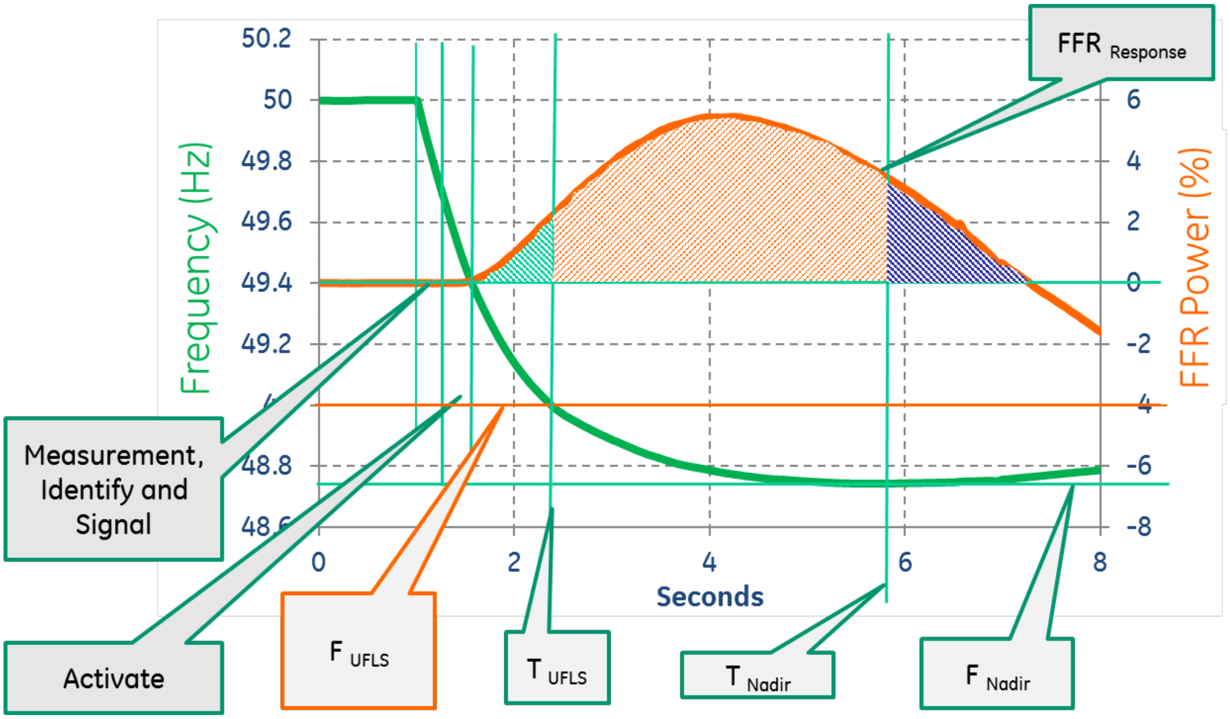

Figure 5 Time Elements of FFR

2.2.2.1 Direct Event Detection

Direct Event Detection with Communicated Response Trigger

If the worst case contingency events are known, and if the type of response required to mitigate the event is also known, then it is possible to directly trigger FFR that rapidly provides arresting energy. These types of schemes are sometimes called Remedial Action Schemes (RAS) or SPS. The basic approach involves:

1. Detect the initiating event (e.g., breaker opening on the AC interconnector, unexpected zeroing of current on the interconnector, etc.)
2. Transmit that information to the FFR device (e.g., HVDC converter terminal or battery storage system)
3. Initiate a pre-programmed action to inject arresting energy into the grid.

This approach bypasses many of the challenges associated with rapid and secure measurement of frequency and/or RoCoF. However, it introduces its own set of challenges, which may or may not be easier to solve. Solutions are application-specific, and depend primarily on grid topologies and communication system performance.

The topology of the problem for South Australia lends itself well to this class of solution. The most severe grid event is well known (non-credible loss of the Heywood AC interconnector). So, the major challenges with this approach involve communications. However, the fact that there are a multiplicity of places where the system can separate means that the design will be relatively complex. Further, as will be discussed at length below, the amount armed for direct actuation must be continuously updated (e.g. by SCADA) for the current operating condition.

The communicated signal is binary, so the message is very short. But the speed requirements often require a dedicated channel from transmitter to receiver. Given the criticality of this function to grid stability, redundant systems would be prudent, and are probably required to meet reliability standards (as is the case in the US). State of the art communications between major transmission substations are typically capable of such functions – as a custom application within the communication network. But given the rapid speed required by AEMO, actual performance capabilities of the South Australia communication systems would need to be evaluated to determine feasibility. Single point-to-point high-speed channels are usually straightforward to implement. Broadcasting signals to several substations may also be possible with similar speed. Issues of cybersecurity must be addressed, and will likely involve encryption.

There are operating installations, using commercially available technology, that appear to meet much of the requirements for the South Australian challenge. For example, there is an operating SPS in California for the Diablo Canyon nuclear power plant[16]. That SPS is designed to detect when two geographically separated circuits both trip, and rapidly initiate tripping of one of the 1GW Diablo Canyon units (so that the other unit will not lose synchronism). Total time for detection and communication is less than 40 ms following the current zero.

Direct Event Detection with Wide-Area Broadcast Options

This concept is similar to the direct detection scheme described above, but rather than transmitting the trigger signal to FFR resources located in major substations, this approach would broadcast the trigger signal to many distributed resources spread over a wide geographic area. The basic approach involves:

1. Detect the initiating event (e.g., breaker opening on the AC tie; and/or current zero)
2. Transmit that information to major transmission and/or distribution substations using the system's high-speed communication capabilities.
3. Broadcast the message from the transmission substations to numerous individual resources (e.g., demand side management (DSM)-enabled loads) using radio or other technology
4. Initiate a pre-armed action that contributes arresting energy into the grid (e.g., rapid load-shedding).

[16] "High-speed control scheme to prevent instability of a large multi-unit power plant", Vahid Madani, E. Taylor, D. Erwin, A, Meklin, Mark Adamiak, IEEE 2007 10.1109/CPRE 2007.359906.

The major challenges for this type of scheme involve speed and security. Given that the triggering signal needs to travel long distances and through several different communication media in series, it may be challenging to meet the speed requirements for this application. Cyber-security is also a challenge, since signal verification schemes would likely be necessary to prevent false actions, and such schemes may introduce additional delay times.

However, with the assumption of well-developed and properly specified fiber optic communications, there are commercial installations in operation which again appear to have the necessary attributes for widespread activation of distributed and heterogeneous FFR resources. For example, there is an operating SPS in Arizona , US, that is sensitive to the loss of 2 of 3 of the nuclear generators at the Palo Verde Nuclear Power Station[17]. (This event is *the* design basis event for frequency response in the Western Interconnection[18] of the US). That SPS is owned by one of the utilities that owns a share of Palo Verde, Salt River Project (SRP). Upon detection of the units' trip, signals to trip load are distributed to many SRP substations. The decision processors have processing time of 1 to 4 ms, detection time for the unit trip is about 1 cycle, and signal latency for communication to all the SRP substations is 7 ms. So, the measure, identify and signal steps (to the substations) is about 2 cycles. That system is redundant and meets current North American Electric Reliability Corporation (NERC) cyber-security rules. These systems can use a variety of protocols (e.g. jungle MUX; MPLS - multiprotocol label switching).

Other more complex (and more expensive) systems can handle a wider variety of processes. There is a new system in operation in Southern California Edison (part of the US Western Interconnection around Los Angeles) that has a "line out" scheme. It detects breaker-open and current-zero in 1 cycle (It is impossible to guarantee that the current will pass through zero, allowing for interruption, faster than 1 cycle). This system uses secure GOOSE multitask routable protocols that are authenticated and encrypted. Again, these meet all NERC cybersecurity requirements. Enabling hardware for that system comes from SISCO[19]. Much of the enabling protocols and practice are in the process of being added to the next version of IEC Standard 61850.

2.2.2.2 RoCoF Detection

The focus of Phase 1 of this project is to examine various FFR technologies to determine how well they can mitigate risks associated with high RoCoF. This requires detection of conditions that require action. This section discusses some of the options and challenges with RoCoF detection for the type of severe frequency event potentially faced by South Australia. Here we address in detail the *Measure* and *Identify* steps discussed in section 2.2.2.

---

17 "Implementation and operational experience of a wide-area special protection scheme on the SRP system", J. Sykes, Mark Adamiak G. Brunello, 2006 Power Systems Conference: Advanced Metering, Protection, Control, Communication and Distributed Resources, 10.1109/PSAMP.2006.285385.

18 In this context, "Interconnection" refers to the synchronously interconnected power system (the western portion of the US and Canada).

19 www.sisconet.com

Measuring Frequency and RoCoF

The immediately obvious solution to the detection of severe events that will require the FFR is to measure the frequency and RoCoF at the device or resource that will respond.

Frequency measurement at a generation resource differs, depending on whether the generator is synchronous or non-synchronous. *All* synchronous generation uses the *speed* of the machine as a proxy for system frequency. That is the origin of the word "governor" – it governs the speed of the machine. There is a close relationship between the speed of a synchronous machine and system frequency, but the two quantities are NOT interchangeable when it comes to controls. In the time frame before speed equals frequency, the inertial response of synchronous machines dominates. In this period, the electrical power out is not equal to the mechanical power in. This inertial behavior is dictated by the physics (Park's equations) of the machine. It is not controllable; nor can it be adjusted. Non-synchronous generation (such as wind and PV) must measure frequency by some other means. This opens the potential for superior frequency response, but it also presents an additional (and new) technical detail than can lead to unintended consequences for overly prescriptive rules.

Therefore, measurement of frequency from observable quantities, i.e. voltage or current, is more of a deduction than a direct measurement. Assuming that measured voltages are sinusoidal[20] (at 50Hz ± Δf), frequency deviation from 50Hz is effectively the rate of change of the angle of the measured phasor. This means that frequency is effectively the derivative of an observed signal, and therefore, it is noisy. That also means that the measurement cannot be made instantaneously, since some filtering is needed to reduce noise. There are a variety of algorithms that have different trade-offs in accuracy and speed, e.g., time between zero crossings, fast Fourier transforms, etc. Control schemes that have substantial system security impact and which depend on rapid response to changes in frequency can require measurement systems specifically built and tested for the application.[21] Complicating matters further is that RoCoF is the derivative of frequency with respect to time, and therefore the measurement is the second derivative of the phasor angle. Because this is a difficult and noisy signal, it requires a significant sample interval to achieve good fidelity.

It is noted that examples discussed in Section 3 are often triggered by *frequency*, not *RoCoF*. Part of the reason for this is that some grids are less concerned with RoCoF and more concerned with frequency. In some grids, RoCoF relays on generators are used to protect them against anti-islanding, and therefore RoCoF is an important metric. In other grids, frequency is the metric used to determine the health of the system and the main concern is avoiding UFLS, so triggering on a frequency threshold is more appropriate. As

---

[20] This assumption good enough for most purposes here, but the sinusoids are distorted and imperfect during severe events, such as before a fault is cleared.

[21] Carlini, EM; Bruno, G; et.al. "Electrochemical energy Storage systems and Ancillary Services: the Italian TSO's Experience", C4-116 CIGRE 2016,

discussed above, an accurate measure of frequency is easier than of RoCoF. However, there are issues with false triggering that are common to both frequency and RoCoF and those are discussed below. Use of RoCoF has potential advantages as well. It is a leading indicator of a frequency disturbance, and can improve dynamic response (not unlike introducing lead to closed-loop control systems).

Phasor Measurements (PMUs) and Wide Area Measurement Systems (WAMS)

PMUs can measure frequency and RoCoF very quickly. For example, Vizimax's PMU can report out 200 frames per second and claim a highly accurate RoCoF calculation. They claim fast settling of RoCoF in the time frame of 1.2 to 3.25 cycles (24 to 65 ms)[22]. Figure 6 and Figure 7 show laboratory test results of their PMU performance. (x axis is time in ms; FE is frequency error, RFE is RoCoF error) If this is as fast and accurate when deployed on South Australia's system, then this may be a good solution, in conjunction with other fast technologies described later in the report, for arresting frequency. Other original equipment manufacturers (OEMs) for PMUs (e.g. GE) also report ability to measure RoCoF in 2 to 3 cycles. However, this is still very much new ground for the industry, and consideration of the ability of these approaches to differentiate local frequency disturbances versus systemic disturbances needs custom evaluation for specific applications.

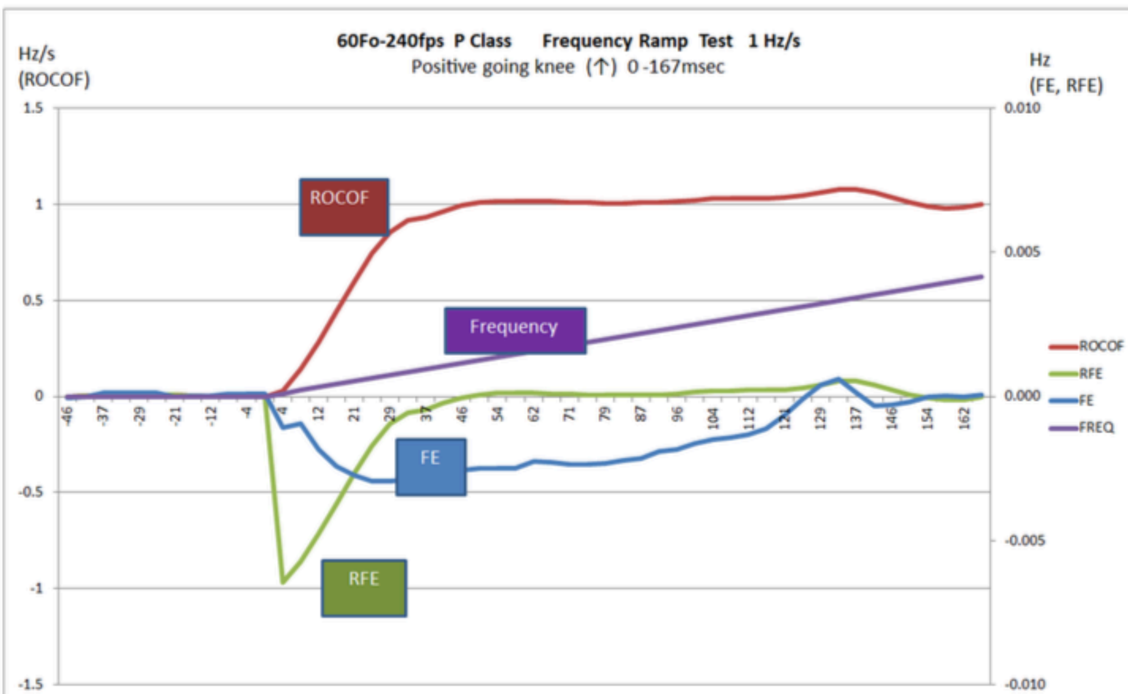

Figure 6 Example of commercial PMU response time and error in detecting frequency and RoCoF at the start of a 1 Hz/s frequency ramp test[23]

22 https://www.vizimax.com/products-services/phasor-measurement-unit

23 Desrochers, A. "Vizimax Energy 3.0 Phasor Meausrement Unit", March 15, 2016.

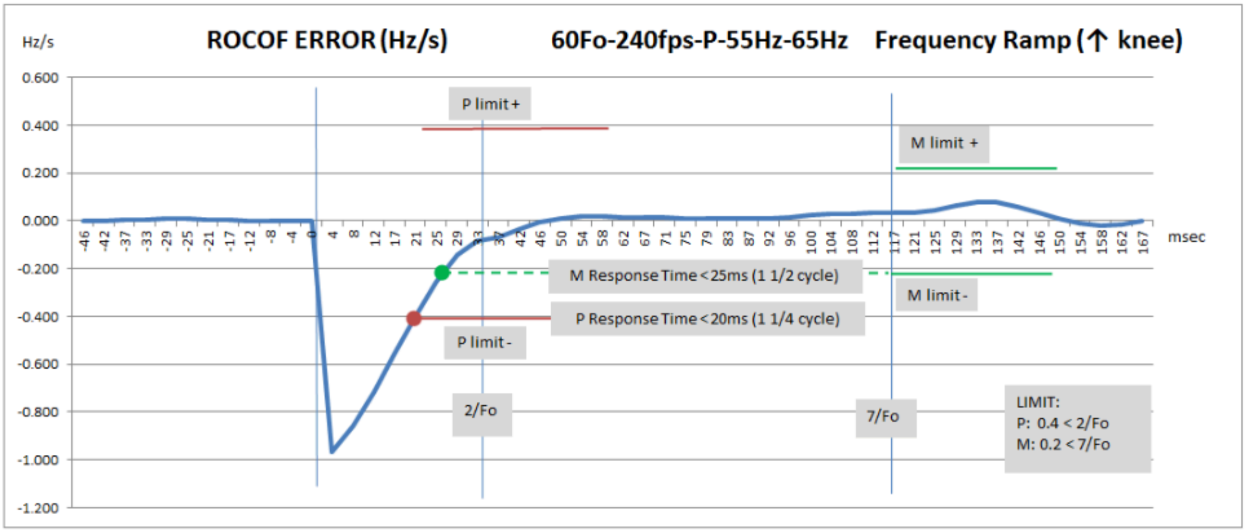

Figure 7 Details of the RoCoF error showing fast settling of RoCoF[24]

PMUs can measure system separation, rather than frequency. Work is being done in several areas on this, including UK National Grid. The intent is to recognize when the system is breaking into pieces, and to use that information to initiate controls that will respond properly to the (sometimes) radically changed conditions that go with the island. This is essentially the problem faced in South Australia. The PMU technology is well established, but the algorithmic customization is still new ground. In these applications, the distributed PMUs require some degree of communication between each other to make good decisions. This takes time. Present objectives are (reportedly) to make these decisions on the order of ½ second. Reaching decisions in less than ¼ second appears challenging, but progress is being made rapidly on these fronts.

Identification Issues and False Triggering

There are two main challenges to fast RoCoF detection while minimizing false triggering:

1. **High RoCoF in one location may not be indicative of severe event** – Assume that a large FFR resource can quickly and accurately measure frequency and RoCoF at its location and that based on this measurement, it identifies a severe event. However, directly after a disturbance, frequency may vary locally. Depending on where this large resource is located and how frequency varies across South Australia, it may overestimate the severity of the event and over-respond or underestimate the event and under-respond. A larger number of smaller FFR resources located across South Australia may mitigate this issue but may be more complex and costly to implement.
2. **Distinguishing between severe frequency events and events that do not require FFR** **–** It is difficult to quickly distinguish between severe frequency events and events that do not require FFR, such as a fault on the tie line that then clears. FFR that is fooled by

---

[24] Desrochers, A. "Vizimax Energy 3.0 Phasor Measurement Unit", March 15, 2016.

disturbances that are not severe under-frequency events may trigger falsely and frequently, unduly impacting frequency. Allowing more time for the RoCoF detection could mitigate this but that is at the expense of FFR response time.

Sampling Window Time

Another critical factor affecting measurement of RoCoF is the period of the measurement. One issue is that a relatively long sampling window may be necessary to distinguish overall grid frequency from local dynamic effects following a disturbance.

For example, Figure 8 shows a 20-degree phase jump event. This is a high, but credible value of loss of one of the two Heywood circuits under heavy loading. Assuming frequency and RoCoF are calculated at the minimum possible timing of every half cycle (10 ms), and that frequency is calculated by measuring time between zero crossings, frequency of the sampling window (red trace, right hand scale) after the event drops to 45 Hz. A 5 Hz drop in 10 ms is -500 Hz/s RoCoF. RoCoF in the next sampling window would be + 500 Hz/s and then zero after that. Triggering FFR based on a measurement of only the first sampling window of 10ms would lead to an adverse response. Longer sampling windows would increase the accuracy of identification of need for FFR. This is just one example of why extremely fast measurement of RoCoF, even if the measurement is *perfectly accurate*, can be problematic.

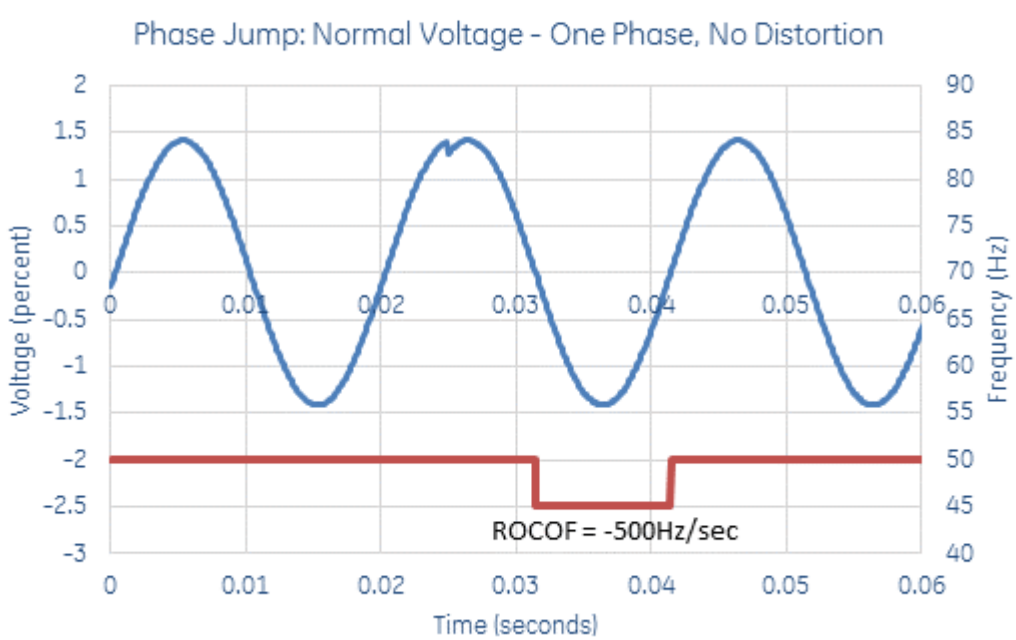

Figure 8 Fast RoCoF sampling window may lead to false triggering

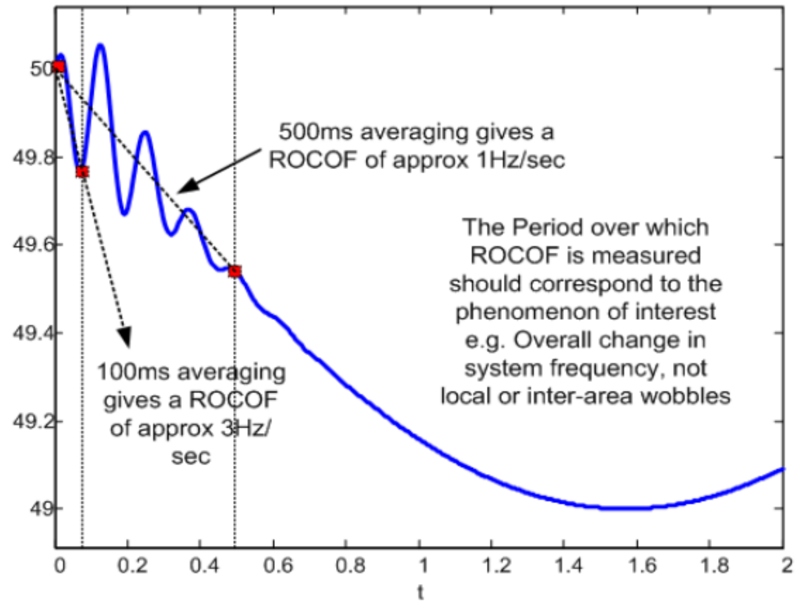

Source: EirGrid and SONI Position Paper, September 2012

Figure 9 Impact of RoCoF Measurement Period on Measured Value

See Figure 9, which appears in a paper prepared by Eirgrid and SONI[25]. In this example, a sampling period of 500 ms gives a 1 Hz/sec RoCoF while a sampling period of 100 ms gives a much higher value that is dominated by a local oscillatory phenomenon and is not indicative of the overall grid frequency. AEMO will need to be cognizant of such effects when attempting to determine whether to trigger FFR based on a rapid measurement of frequency or RoCoF. A separate issue, which is the subject of a separate AEMO effort, concerns the sensitivity of equipment to RoCoF. Unlike the *deliberate* response to measured RoCoF, that is the subject of this investigation, if equipment (e.g. generation) trips *because* of the high RoCoF in an undesired (and perhaps unexpected) fashion, this high local RoCoF could presage systemic problems – including risk of cascading failure.

Frequency measuring devices that can meet the highly precise and fast requirements of power electronic control during and immediately following system disturbances are often specially created for the specific technology. For example, GE developed a Continuous Wave Frequency Transducer (CWFT) for use in HVDC modulation functions and turbine-generator power system stabilizers (PSS). This device uses phase-locked loop methods to minimize sensitivity to non-fundamental frequency distortions in the input signal waveforms. This transducer is directly integrated into these products to achieve fast and stable performance in power swing damping applications. This type of transducer could possibly be adapted for other types of applications.

25 RoCoF Modification Proposal–TSOs' Recommendations, http://www.soni.ltd.uk/media/documents/Archive/RoCoF%20Modification%20Proposal%20TSOs%20Opinion.pdf

If RoCoF measurement and event detection is used as a main method for triggering FFR, it would be wise for AEMO to study their frequency response from real events to determine the size of the sampling window to manage the tradeoff between accuracy in event detection and time to respond.

Local Measurement of Frequency and RoCoF

Examples of events usually show system frequency to be a uniform quantity. However, in the initial period following a large disturbance, system dynamics result in multi-modal swings. Consequently, until these inter-area swings damp out, frequency varies with location. This can be observed in Figure 10. The significance of this locational effect for mandating and paying for frequency response, especially FFR, can be significant. Consider the traces in Figure 10. The disturbance occurs at 1.0 seconds in the simulation. The buses near the lost generation (teal blue) drop very rapidly (i.e. high initial RoCoF), but the bus farthest from the disturbance, the orange trace here (which is more than 1000 miles away from the disturbance), doesn't "see" the event for nearly 3 seconds. A further example, with time synchronized phasor measurements from a large loss-of-generation event on the edge of the US Eastern Interconnection is shown in Figure 11. The farthest bus shown does not see the event for about 4 seconds. It is of interest to note that the sign of the frequency at this remote bus is the opposite of that on the buses nearest to the initiating event after 3 or 4 seconds. Another excellent example of this effect for the Great Britain grid is shown in a very recent CIGRE paper[26].

These examples are for systems that are electrically and geographically larger than South Australia. Consequently, the time for the event to propagate across these example systems is longer than would be expected in South Australia, nevertheless, this information demonstrates that FFR compliance based on when the event occurred, rather than when the bus could observe the event, would have very different conclusions.

A discussed above, AEMO has performed analysis that shows a range of possible RoCoF that is dependent on operating conditions, both present and possible in the future. Measuring RoCoF and frequency with a view towards triggering controls for arresting frequency collapse at the highest rates presently identified by AEMO as possible, i.e. exceeding 4 Hz/sec, is largely uncharted territory. At the very least, detailed design and evaluation studies would be required to qualify specific detection algorithms. There is some precedence with bespoke technologies, such as those that have been developed for individual HVDC projects and energy storage projects, but extending these concepts to widespread deployment on a multiplicity of installations of different technology, vintage and manufacturer, is, at present, unprecedented.

[26] See Figure 6 in the paper, Advances in Wide Area Monitoring and Control to address Emerging Requirements related to Inertia, Stability and Power Transfer in the GB Power System, D.H. Wilson et al, CIGRE Paper C2-208, Paris, August 2016.

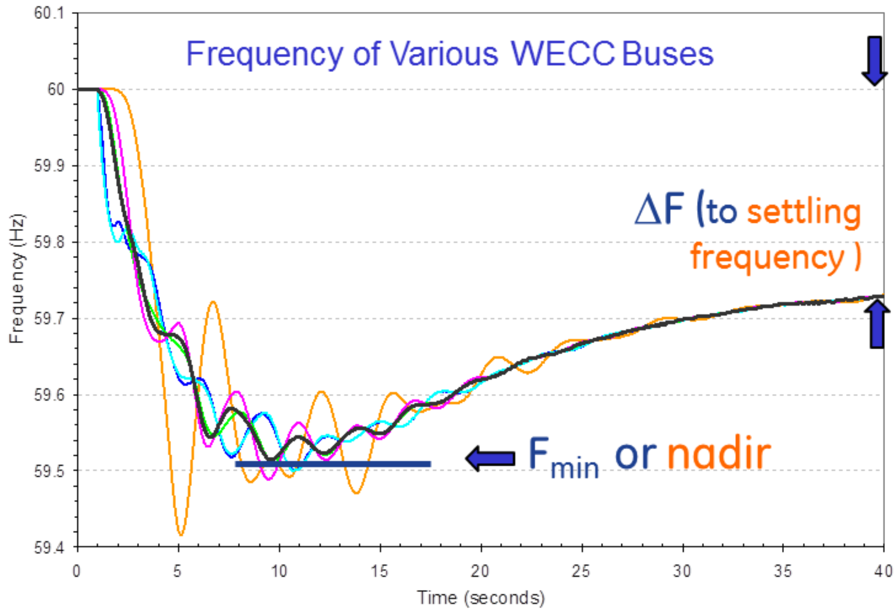

Figure 10 Bus Frequencies for Large Disturbance in USA Western Grid

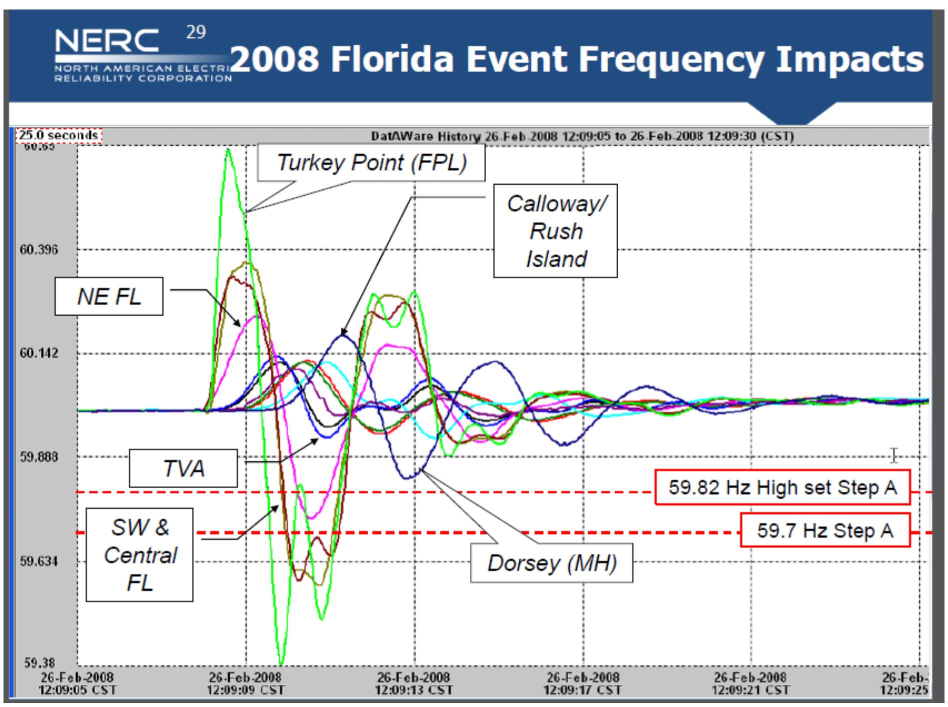

Figure 11 Geographic Distribution of Bus Frequency for 2008 Florida, US Event[27]

---

27 Robert Cummings, NERC, 2012 "Florida Disturbance Feb 26, 2008".

Example of Challenge of RoCoF and Frequency Measurement Differentiation

Making extremely fast differentiation between critical events (for which FFR is required) and other events can be challenging. Here a very simple example is presented. In Figure 12, three frequency traces are shown. The blue trace is a measurement provided by AEMO of the November 1, 2015 separation event, and the red trace a simulation of the event, performed on a greatly simplified model created specifically for this work. (The model captures some of the essential elements of the South Australia system as it relates to the separation and RoCoF issues, but is in no way intended to be a substitute for the detailed representations maintained and used by AEMO for analysis. The model and more results are presented and discussed at length in Section 6). This event is representative of the sort of event for which FFR would be desirable. The RoCoF in this case is less than 1 Hz/sec. The third trace is for simulation of a line fault and clear event, for which FFR would not be necessary. Clearly, the green event represents a hugely different event.

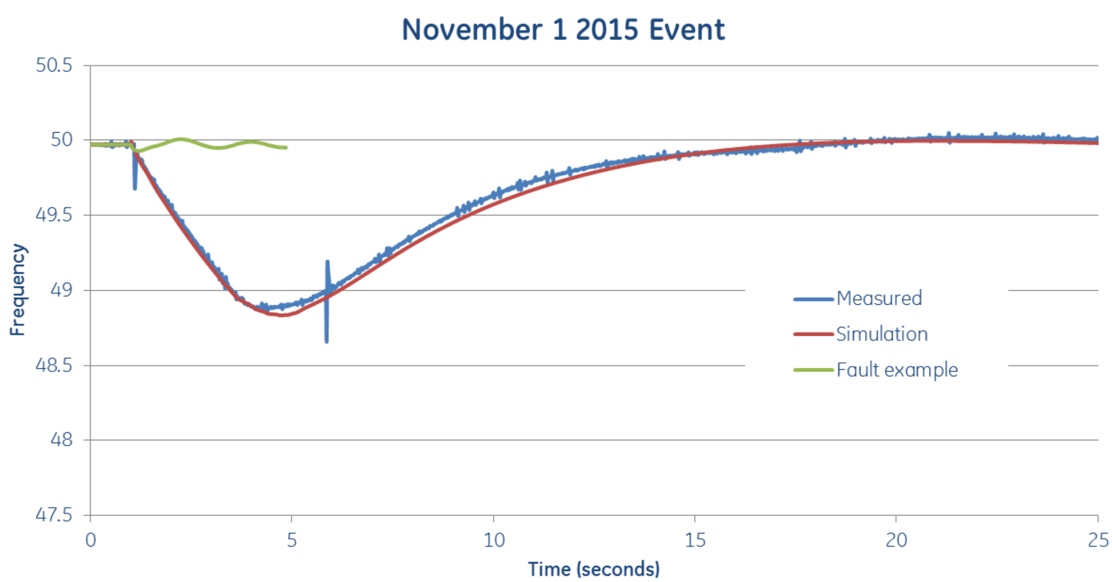

Figure 12 November 1, 2015 Event (Measurement and Simulation) with Simulated Fault

But for this discussion, we are concerned with very rapid detection and response to the event. So, in Figure 13 the same information is plotted, but zoomed in for the first 0.3s of the event. Whatever was used to provide the AEMO frequency measurement (blue trace) took about 80ms to respond, then it overshot, finally giving reasonable frequency information after about 0.2s. This particular device could not have provided useful RoCoF information faster than 1/5 second. The frequency traces are from the computer simulation. The simulation uses positive sequence representation (as do essentially all stability programs), and consequently frequency calculations are subject to the limitations of sequence calculations and network solution methods. It is not obvious that the initial jump up in frequency in the separation simulation is physically meaningful. The two simulated events also have essentially identical frequency deviation (about -0.070Hz –

blue arrow) after 0.1s (red arrow). In any event, for this admittedly very simple example, differentiation between these two dramatically different grid events in a time period less than 100ms would be problematic at best. The point here is that, even with very high fidelity calculation of frequency, dependence on purely local signals to rapidly detect events may be subject to false triggers.

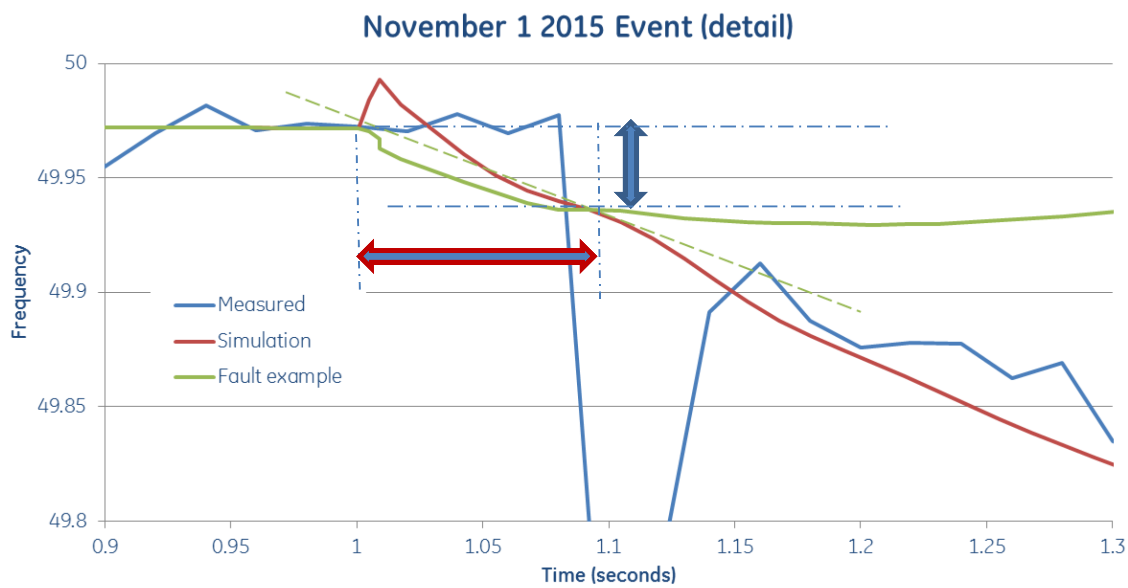

Figure 13 Detail - November 1, 2015 Event

Emerging Microgrid Technologies

There is an enormous amount of research and development under the umbrella of "microgrids". The technical challenges faced by South Australia have a high level of commonality with the objectives and constraints of (at least some definitions of) microgrids, and indeed one could argue that South Australia *is* (or becomes) a microgrid when it separates from the rest of the NEM. Both the physical and MW scale of the problem is large (compared to the subject of most work), but the issues of maintaining frequency and stability with a predominance of inverter based resources and low inertia is essentially the same. The proportional **scale** of the disturbances for which South Australia is trying to maintain stability is nearly unprecedented however.

Virtual Synchronous Machines

The challenges introduced by displacement of synchronous generation by inverter-based resources gives rise to a relatively simple question (with complex answers): "Why not make the devices behave like synchronous machines?" And indeed, there is an emerging consensus that inverter-based resources may need to emulate the electrical characteristics of synchronous machines in systems that are predominantly or solely supplied by them. This is not a new concept (see 1995 battery example in section 3.3), but has costs and complexities.

The issue is discussed further in section 4.5, but at present such controls are not (to our knowledge) available in commercial wind and PV solar generation.

Other System Risks

The focus of this report is RoCoF for the NEM, with attention to the challenges associated with separation of South Australia from the rest of the NEM. Given recent events[28], it is worth reinforcing that technologies introduced for managing high RoCoF should, ideally contribute to providing other stability benefits, particularly helping to avoid separation of South Australia from the rest of the NEM.

2.3 Complexity, Costs and Robustness

There are tradeoffs between the various approaches presented above. Cost, complexity, speed, and risk of false triggering are all inter-related. Discussion above mentioned the direct event detection versus local RoCoF detection methods. For example, detection that the AC link is down is relatively inexpensive but costly communications links are required to transmit that information to the FFR resources. In contrast, local RoCoF detection at the FFR resources eliminates the cost of communication but at the expense of installing multiple RoCoF detection devices.

Similarly, there is a tradeoff between the number of devices and the risk of false triggering. Because local frequency may not be a good indication of a severe event, multiple RoCoF detection devices across South Australia (associated with their own FFR resources or with communications links to FFR resources) can reduce the risk of false triggering by giving a better indication of system frequency health. This risk reduction comes at the expense of higher cost in RoCoF detection devices or communications links.

And then finally, there is the risk of false triggering due to the difficulty of discerning an event very quickly from the frequency measurements. The tradeoff in detection time and risk of triggering when response is not needed must be considered. Because events that do not require FFR can initially look like events that do require FFR, it may be better to give more time for detection, as EirGrid has done with their proposed 500 ms averaging window. However, this squeezes response time. For example, 500 ms would be too slow for AEMO to respond to a 4 Hz/s RoCoF. There is not a specific "right" answer to this trade-off (e.g. "OK up to X Hz/sec RoCoF"), as the figure implies. More quantitative context on the requirement for speed of response is provided in Sections 5 and 6.

To design a robust solution, a detailed examination of the tradeoffs discussed here would be wise. Further work needs to go into:

- Understanding the impacts of false triggering
- Understanding how easily severe frequency events in South Australia can be discerned from other disturbances in South Australia

28 AEMO "Update report – Black system Event in South Australia on 28 September 2016" 19 October 2016.

- Understanding how local frequency varies across South Australia due to a severe event
- Testing how well commercial solutions can detect severe events in South Australia
- Quantifying the costs of various options

2.4 Building Confidence in What is Possible

In summary, AEMO's needs for the South Australia grid for the most extreme events are at the boundaries of existing technical capabilities and beyond industry experience. AEMO's overall speed of response targets are significantly shorter than those of others facing similar issues (Great Britain, Ireland, etc.). Speed requirements on individual components that might be part of an overall solution are also faster than previous applications, and may not be technically possible with adequate resolution and security.

In the process of narrowing the possibilities and developing technically feasible conceptual designs, laboratory testing of the essential components and subsystems is critical to success. Performance characteristics of transducers, communication systems, and control systems are application dependent, and cannot be derived from published specification sheets when requiring performance at the limits of their capabilities. There is an axiom that seems appropriate here; "One test is worth more than a thousand expert opinions."

3 TECHNOLOGY CAPABILITIES FOR FFR SERVICE

This section describes technologies that may be able to provide a useful FFR and their capabilities and characteristics.

3.1 Overview

There are many options for technologies that may be able to provide (the actuation portion of) FFR. Because conventional synchronous technologies, (e.g. adding synchronous capacity or adding new interconnectors) are well-understood, they are not included here. Instead, the focus of this report is on inverter-based technologies and other technologies that may not provide an inherent inertial response. The analysis is presented by technology in the following sections. Ultimately, this is input to a broader evaluation of options for AEMO, that includes the options identified in this report along with "conventional" alternatives.

Delivery of FFR services by the various technologies will depend on the physics of the technology and on control. The distinction is important, and, is reflected in our discussion. In so far as information is available, the capability of currently commercially available controls is noted and described. However, it is important to recognize that provision of specific services by technology is often catalyzed by market demands: even if no OEM has offered a functionality, that capability might be offered if the market demanded it.

The discussion below also tries to provide insight into what is possible subject to limitations of practical controls. Consideration of stability, latency, efficacy and robustness are included. For context, it is of note that NERC has discouraged unduly fast response from inverter-based resources in various instances due to the risk of unstable or interacting controls. "Just because you can make it fast, doesn't mean you should". Good engineering practice and reliance on good judgment is essential. It is challenging to make prescriptive rules, particularly market rules, that will solicit and reinforce behavior by participants that best serves the needs of a system. We have tried here to provide physical insight that will aid in making good engineering decisions.

It is important to note that the problems facing South Australia are extreme. The phenomena of interest are really fast. Insights as to what technologies will be successful won't come solely from equipment specification sheets. Further, equipment that can act rapidly under ideal grid conditions, may not perform as well under weak grid conditions. The discussion provided below, while created with the situation in South Australia as a significant consideration, is intended to provide more general insight, applicable to the entire NEM. Issues of economically maintaining good frequency response, with consideration of all available technologies to do so, is of interest for the entire country.

3.2 Wind Turbines

3.2.1 Overview of the Physics of Inertia-based FFR from Wind Turbines

Most modern large-scale (>1 MW) wind turbine generators (WTGs) rely on sophisticated power electronics to maximize wind power production, control turbine speeds and to provide a range of performance functions. In these variable speed machines, the power electronic enabled controls improve energy capture and provide characteristics that are beneficial to the grid. The economics and physics of wind generation demand that the speed of rotation of the blades must be adjustable with different wind speeds. The industry has almost unanimously settled on the use of power electronics to allow the generator speed to vary, and this electrically decouples the generator speed from the grid frequency. Consequently, unlike simpler induction generator based systems, turbine-generators with these controls do not naturally provide inertial response.

The existing population of WTGs in South Australia includes a mixture of different vintage and technology wind turbines. Looking forward, essentially all new utility scale wind generation being built in the world uses electronically enabled variable speed generation (i.e. Type 3 and 4).

In the case of induction machines and the truly synchronous machines (very rare "type 5" for wind turbines), there is a direct connection between the power system and the machine. When there is frequency decay on the power system, the induction machine will increase its output temporarily because of the slip change. The induction machines make a limited contribution to system inertia while the truly synchronous machines will inherently add inertia to the system the same way a hydro or thermal turbine would[29]. Type 2 wind turbines, of which there are some still in operation in the NEM, are induction machines with controlled resistance on the rotor. This allows some degree of power control during swings. This is discussed separately below in Section 3.2.3.3.

The basic design of the present class of converter based technology (Type 3 – aka "double-fed – DFG" or "DFIG" generators) and 4 – aka "full converter generators"), however, does not include any inertial response unless explicitly designed to do so. Both the DFG and full converter generators employ a back-to-back converter to connect to the power system, enabling speed variability for the turbine rotor.

Unlike conventional synchronous or induction generators, the delivery of active power from variable speed wind generators is almost entirely controlled by the power electronics. The power electronic controls allow nearly instantaneous adjustment of electrical torque on the generator, and therefore power delivery, which is essentially independent of the terminal bus voltage angle, and rate of change of angle – i.e. frequency. Controlled inertial response is possible because of this extremely fast response. This is a fast, controlled response, rather

[29] Kundar, P., Power System Stability and Control, 1994, McGraw-Hill, Inc., New York, ISBN 0-07-035958-X.

than the inherent, uncontrollable response of all synchronous machines. This response is transiently decoupled from the mechanical angle of the wind generator rotor. Therefore, it is possible, though not necessarily desirable, to more-or-less instantly change (a few cycles) electrical power delivery. As discussed above, in the context of the challenge facing South Australia, this few cycles is the "activate" time. It is worth reiterating the note above (in Section 2.2) that inverter-based devices, including wind turbines, could be designed to behave as virtual synchronous machines. This presents a possible, even likely, path forward for the industry faced with systems evolving towards occasional operation with zero synchronous inertia. This is not the present offering, and this evolution comes with costs and challenges, as discussed further in Section 4.5.

3.2.2 A Note on Terminology

The language of the industry is confusing and unsettled. Early work talked about "replacing" the lost synchronous inertia from displaced thermal machines. Further, the available controls generally are *enabled* by accessing the energy stored in the inertia of the turbine. So, in that sense, these are inertial controls. BUT, the physics of the behavior is fundamentally different: the power electronic devices DO NOT mimic Park's equations…i.e. they don't "look" like synchronous machines. The power injected to the grid is one variety of FFR. Unlike some of the other technologies discussed in subsequent sections, the energy available is relatively limited, so the response cannot be sustained throughout the time period when secondary frequency response returns system frequency to normal.

The discussion at times is becoming heated: some people at the US Federal Energy Regulatory Commission strongly dislike the term "synthetic inertia", but they also dislike simply calling it FFR, since the inertially enabled control cannot be sustained. In contrast, what is generally called PFR control from wind turbines is enabled by both the pitch control and the power electronic torque control (discussed further below).

Through-out the balance of this discussion we will use the notation "**inertia-based FFR**" to capture this specific behavior.

3.2.3 More Discussion of Inertia-based FFR Controls for Wind Turbines

The power delivery of the wind turbine-generator is limited not only by the available wind, but by the physical limitations of the components of the WTG. Most critical are aero-mechanical ratings and speed limits. The lift of wind turbine blades is a strong function of blade speed relative to wind speed. These well understood relationships[30] are reflected in the control algorithms of modern, variable speed WTGs. Blade speed and rotational speed are directly proportional, so turbine controllers target optimal rotational speed for power production. The speed of the turbine is dictated by the mechanical torque delivered from the blades and the electrical torque removed by the generator. When in balance, the turbine speed is constant.

[30] Ackermann, Thomas; Wind Power in Power Systems; Wiley; Sussex, UK, c. 2005; pg 25-50.

Inertia-based FFR Controls for wind turbines are based on temporarily making the electric power delivered to the grid exceed the mechanical power being captured from the wind. The source of energy for this extra power is the stored rotational energy in the turbine rotor and drive-train. To extract that energy, the rotor must slow down.

A key point is that slowing the turbine tends to reduce the aerodynamic lift, thereby reducing the delivered mechanical shaft torque and exacerbating the speed decline caused by increased generator electrical torque. This positive feedback tends to push the blade towards aerodynamic stall, which must be avoided. The inertial control must provide margin above stall, and is consequently limited when the initial rotor speed is low. This means that the energy available response is limited whenever the wind speed is at or below rated. Further, the amount of energy available drops with wind speed (and therefore turbine power). At low wind turbine power levels, the available energy of the inertial response starts to decline rapidly below about 50% rated power, dropping to zero below about 20%.

Inertial energy extracted by slowing the rotation of the turbine must ultimately be recovered. After the initial increase in electrical power, it must temporarily drop below the mechanical power to allow the energy to recover, reaccelerating the rotor. The amount of energy that needs to be recovered therefore includes (a) the extra arresting energy delivered to the grid and (b) the energy NOT collected during the turbine speed drop because of reduced lift. One can think of this second term as interest to be paid on the energy borrowed from the rotor. This can be a critical consideration in systems with high penetration of wind PLUS limitations on primary frequency response. Ultimately, inertia-based FFR from wind turbines buys time for primary frequency response to act. This is a key fundamental observation: for operating conditions at or below rated wind speed, *the only NET value provided to the system from inertia-based FFR essentially ends with the frequency nadir. That is, the value is in energy delivered during the arresting period.* Withdrawal of inertia-based FFR should, ideally, begin at the nadir and be coordinated with the rate of rise of the primary frequency response during the post-nadir recovery period. The coordination of withdrawal can be an important consideration, if done too quickly relative to the rise in power injection from primary frequency response, the frequency will fall again. But, done too slowly, the turbines will incur larger than necessary rotational energy deficit that must be recovered.

Another aspect of the aero-mechanical rating of the wind turbine is maximum mechanical loading. When the wind speed is above rated, the turbine controls reduce lift by decreasing the angle of attack through pitch control. Under these conditions, the available wind power is greater than turbine rated power. It is possible to increase the captured wind power, using pitch control, to temporarily exceed the steady-state rating of the turbine. Thus, at higher wind speeds, inertia-based FFR can transiently provide the extra power for inertial response through pitch control - increasing wind power to the rotor, rather than just extracting stored inertial energy by slowing the rotor. Under these conditions, the decline in rotor speed is less and the energy recovery is minimal or non-existent.

The limitations of the electrical system must also be respected. The electrical rating of the converter and generator are matched with the wind turbine. As with most electrical machinery, the equipment has short-term capability that exceeds the continuous rating.

There are, nevertheless, current and voltage maxima that must be respected. The inertia-based FFR control must respect these limits and coordinate with all other controls that can drive the electrical loading on the converter and generator.

It can be a significant advantage that, unlike the inherent response of synchronous machines, inertial WTG response is dependent on active controls and can be tailored, within limits, to the needs of the power system. However, the response is shared with controlled variations in active power necessary to manage the turbine speed and mechanical stresses. These stress management controls take priority over inertial control. Turbulence may mask the response for individual turbines at any instant in time, but overall plant response will be additive. It is an important point for this discussion to note that in broad terms these constraints apply to all modern variable speed wind turbines, regardless of manufacture. But since the primary mechanism to deliver arresting power is via control, there is wide latitude available to the equipment designers to implement different controls. In the following subsections, we will discuss the two general classes of controls, which we have termed "closed-loop and open-loop".

3.2.3.1 Inertia-based FFR – Closed-loop Control Philosophy

Closed-loop inertia-based FFR features have been available commercially for many years[31]. In this context, "closed-loop" means that the control action (i.e. active power injection to the grid) is continuously varying based on continuous measurements (of frequency). Some offerings were developed in response to the 2006 Hydro-Québec grid code[32] to meet the equivalent energy (kW-sec) contribution of a synchronous machine with 3.5 MW-sec/MW inertia for the initial 10 seconds. Given the systemic needs, and the Hydro-Québec requirement, the overall control that GE designed was intended to provide similar arresting energy to that of a synchronous machine. A further objective of the control was to use the finite energy available (i.e. as constrained by avoiding stall and respecting equipment limits) to **have the most beneficial impact on improving the frequency nadir.** The behavior of the control was therefore tailored to work well with a system that (a) has adequate, but relatively slow, primary frequency response, and (b) reaches its frequency nadir in several seconds following a loss-of-infeed event. The net result is a control that is rather different from the inherent response of a synchronous machine.

Part of inertial control design constraints are to meet performance objective over most of the turbine operating range. The specified inertia constant included in the Hydro-Québec grid code is representative of large thermal generation. This target is met, for instance, when the wind turbine dynamically varies the real power by about 5% for 10 seconds in response to a large, short-duration frequency deviation on the power system. Thus, the energy delivered is

[31] Nicholas Miller, Kara Clark, Reigh Walling, "WINDINERTIA™: CONTROLLED INERTIAL RESPONSE FROM GE WIND TURBINE GENERATORS", MIPSYCON 09, October 2009, Minneapolis.

[32] "Transmission provider technical requirements for the connection of power plants to the Hydro Quebec Transmission System", March 2006.

on the order of 0.3 MW-sec/MW when the turbine is operating near rated. Remember, this is NOT the inertia (which has the same units), but the arresting energy that is delivered.

Hydro-Québec was the first transmission owner to require wind plants to contribute to frequency regulation by using the inertial response. In 2010, Hydro-Québec started integrating the first wind plants equipped with this feature in its network. Others in Canada have followed, and several places in the US (notably ERCOT) are considering either requirements or qualification of inertia-based FFR for inclusion in an FFR ancillary service market.

The fact that some of these original designs were primarily based on the requirements for Hydro-Québec has some relevance to Tasmania, as both systems have dynamics dominated by the transient response of hydro generation. The performance discussion and examples in the following section are based on a simplified representation of HQ (at low load and what at the time was considered high wind).

A simple representation of the critical elements is shown in Figure 14. Frequency error is simply the deviation from nominal. A positive frequency error means the frequency is low and extra power is needed. The deadband suppresses response of the controller until the error exceeds a threshold. This limits the inertia-based response to large events. The continuous small perturbations in frequency that characterize normal grid operation are not passed through to the controller. The frequency is extracted from the wind turbine phase-locked-loop (PLL). This frequency and synchronizing signal is also used as input to the current control. The speed of the frequency transducer relative to the speed of the control *as originally developed*, is very fast – on the order of 50ms or less.

The Power Shaping block produces a driving power response signal. The power shaping block includes a "lead" term; this is, in a sense, sensitive to RoCoF. This signal is further filtered and coordinated with other WTG controls, particularly the generator power order in the Power Coordination block. Finally, the inertial response signal is limited, sent to the turbine electrical control which adds to the generator power order. This net command is implemented by the WTG converter controls, ultimately resulting in power delivered to the WTG terminals.

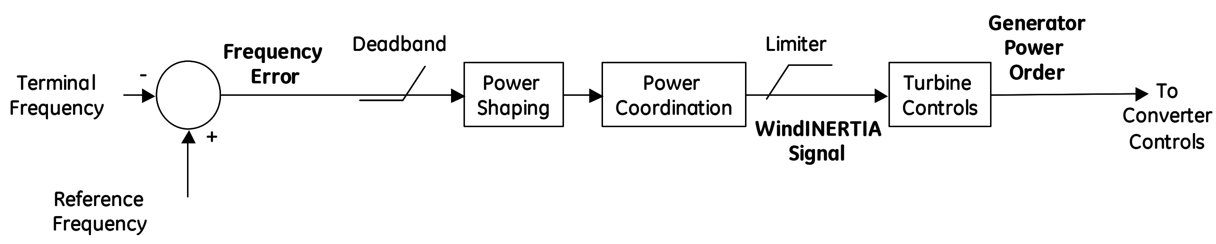

Figure 14 Functional Representation of a Closed-Loop Inertia-based FFR Control[33]

---

[33] Nicholas W. Miller, Kara Clark, Robert W. Delmerico, Mark E. Cardinal, "WindINERTIA™: Controlled Inertial Response from GE Wind Turbine Generators", CanWEA, Vancouver, BC, October 20, 2008.

There are several differences between this inertia-based FFR response, and the inherent inertial response of a synchronous machine. First, and most important, the control is asymmetric: Inertia-based FFR (as offered by OEMs today) only responds to low frequencies. High frequency controls are handled separately, by a different controller that can, if necessary, provide sustained response, as discussed in Section 3.2.4.1. Second, the deadband ensures that the controller only responds to large events – those for which inertial response is important to maintain grid stability, and for which seriously disruptive consequences, like UFLS, may result. It is important to note here that these two differences are not fundamental physical requirements for inertia-based FFR, but rather a result of the design objectives (i.e., help the grid for big loss of generation events).

Finally, an inertial-based FFR response means the speed of response is a function of the control parameters. In the example shown, the response was tuned to provide good coordination not only with inertial response of other generation on the system, but with governor response of conventional generation as well. The ability to tune inertial response (including shutting it off) provides the planning engineer with an additional tool to manage system stability.

Field test results of this inertial control for various wind speeds on a single wind turbine are shown in Figure 15. These tests were run around 2008. The field data was generated by repeated application of a frequency test signal to the control. The results, at various wind speeds (as listed in the legend), were then averaged and plotted. These are tests of the IBFFR only, so in each test, the turbines are producing the maximum power available for the wind condition (i.e. they are not curtailed). Below rated wind speed (<14m/s) the results clearly demonstrate the inertial response and recovery. Above rated wind speed the inertial response is sustained by extracting additional power from the available wind (i.e. short-term overload of the WTG). [These tests were on type-3 wind turbines, but the distinction between type-3 and type-4 is largely irrelevant for this function. The difference in implementation between OEMs is more important.]

These cases all have an initial rate of power rise of about 3.3%/sec (meaning % of rating, not actual power; this works out to 50kw/sec for this individual turbine). At the time this work was done, the upper design limit for rate-of-rise of power was 20%/sec. So, for this (AEMO) discussion, the gains could be increased to have that rise time. Other OEMs have claimed possible rise times up to 30%/sec. Since inertial response *for this design* is limited to 10% of nameplate, the turbine can reach full response in 0.5 seconds, once the control is activated by crossing the frequency deadband.

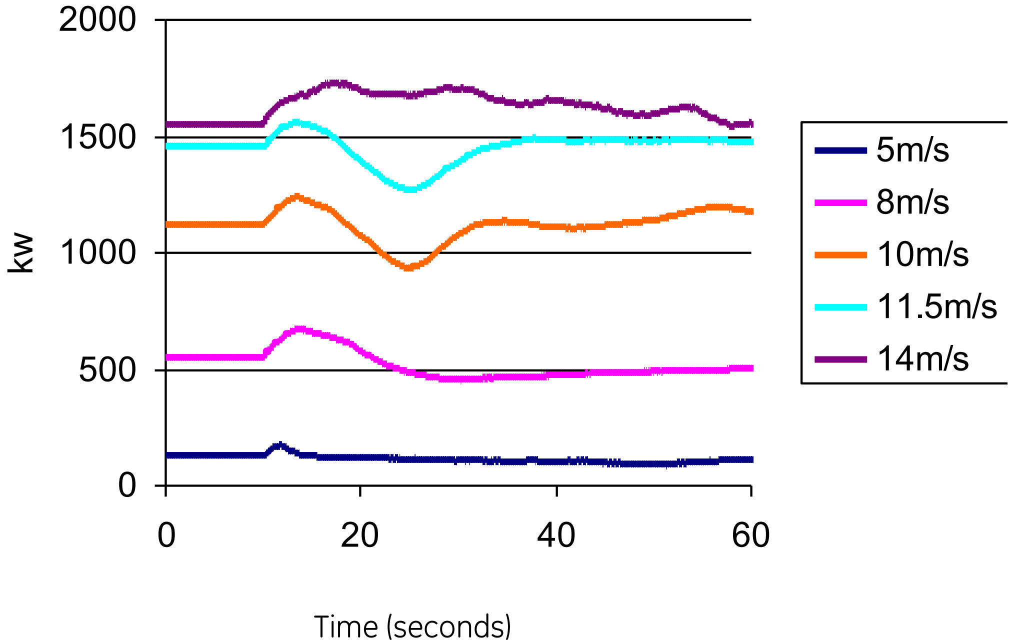

Figure 15 Field demonstration of one inertia-based FFR response [23]

Similar results from Vestas are shown in simulations of Figure 16. For the event, the figure shows a range of responses, as is typical for wind turbines running with actual wind conditions. The response of turbines is of the order of 7 or 8%, with about 5% of the rise occurring over a period of less than one second. The variability in response is an important reality of individual wind turbines. Dictating exact behavior of a plant, and certainly of individual turbines, is not likely to be possible with any available wind turbine control technology. It is also of note that the Vestas authors are dismissive of RoCoF sensitive controls, regarding them as insufficiently robust: "…makes this type of control a non-viable candidate".

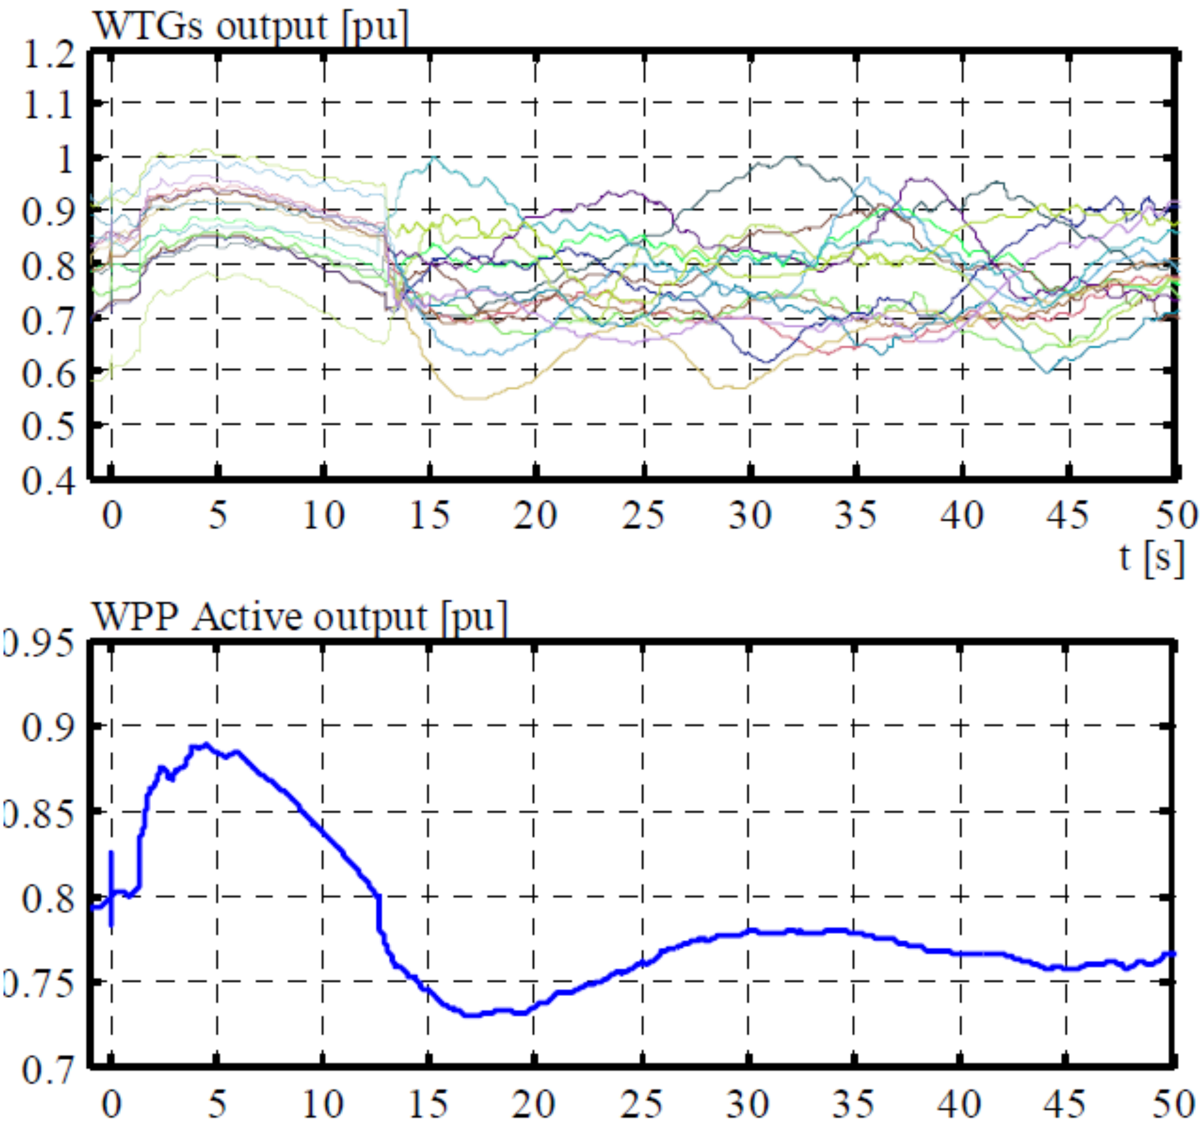

Figure 16 Wind Inertial response (simulations)[34]

The diversity of response is shown in the sequence of traces of Figure 17. The simulated traces are from the control work by RePower (now Senvion) performed in conjunction with Hydro Quebec.[35] Note that all WTGs are subject to the energy recovery period, which for these examples occurs after about 10 seconds of response. The various control schemes employ slightly different schemes, but all authors point out the risks associated with excessive

34 Peter W. Christensen, German Tarnowski, "Inertia for Wind Power Plants – State of the Art Review 2011", 12th Wind Integration Workshop, 2011.

35 M. Dernbach, D. Bagusche, S. Schrader, "Frequency Control in Quebec with DFIG Wind Turbines", 12th Wind Integration Workshop, 2011.

recovery power. The following discussion provides additional context to this important consideration.

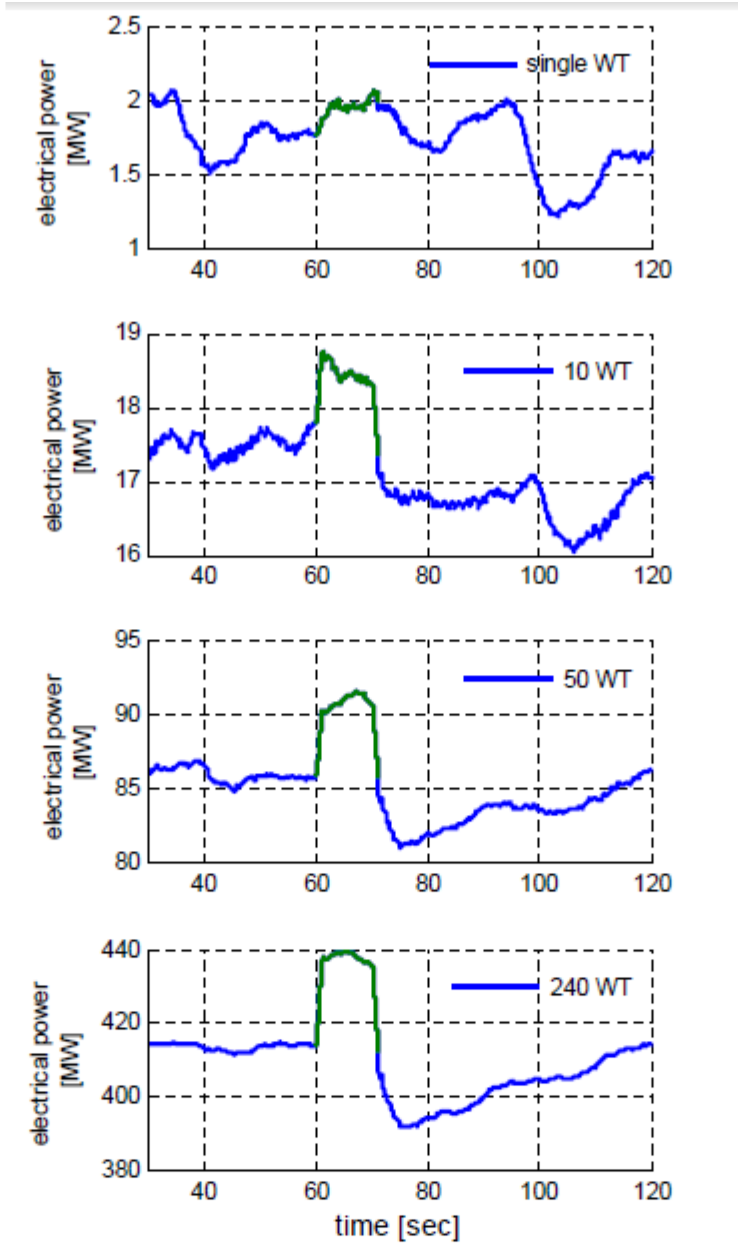

Figure 17 Sequence of Aggregate Response of an inertia-based FFR[35]

3.2.3.2 Inertia-based FFR – Open-loop, triggered Control Philosophy

The need for extremely fast rates of response for South Australia may drive a requirement for an approach that provides a predetermined response to a triggering signal. As discussed above, the trigger could be based on local information such as measured RoCoF or frequency, or from a communicated triggering signal from elsewhere. (The considerations outlined above for measurement versus direct communication apply here). The key distinction is that the response, or more specifically *at least the initial response,* is predetermined once the trigger occurs. There is precedence for this type of control, since it is a form of SPS. In the case of South Australia, making closed loop control fast enough for the trip of the AC interconnector raises concerns over control stability and unintended consequences. Nevertheless, open loop responses are not without drawbacks as well. There can be (institutionally) a tendency to focus designs on worst case scenarios, which will drive very aggressive behavior. This can result in unintended, and undesirable, outcomes for other normally less interesting events. Risks of overshoot, unnecessary stress (including torsional duty not only on the turbine, but on other rotating equipment on the grid) must all be considered. In any event, the analytical burden of either option is not trivial.

One type of open-loop control is shown in Figure 18. This figure is the PFR requirement by Ontario's Independent Electricity System Operator (IESO), recently added in June, 2016[36]. The orange annotations were added by us. The figure represents two "cases" of response. The idea is to apply a minimum performance requirement for wind plants. "Case 1" is for a frequency event that is sustained for more than 10 seconds. The basic idea is that the WTG increases its electric power by 10% over the power production level that immediately preceded the event. The turbine power is required to ramp up to at least 110% of this P_o within 1 second, and sustain it for 10 seconds before being allowed to ramp down *below* P_o. In "Case 2", the control is allowed to revert to the recovery phase when the frequency returns to within a small deadband of nominal. The areas on the figure represent, respectively, the extra energy delivered by control (Area 1) and the recovery energy (Area 2) not delivered to the grid to allow the WTG to return to normal speed. The amount that power drops below P_o in the recovery phase is noted as ΔPr. The IESO notation suggests that this be no more than 10%.

The "area 1" is almost, *but not quite,* the arresting energy delivered by the control. The difference is that the "release" of the control is based on frequency returning to normal, and not on the time of the frequency nadir (which coincides with RoCoF passing through zero and becoming positive. The figure carries some serious risks for practical application. Specifically, the relationships between "Area 1", "Area 2" and "ΔPr" are emphatically *not* independent. The physics of the problem dictate that the longer the turbine runs below its optimum speed, the more recovery energy is needed. Thus, as ΔPr shrinks, Area 2 increases, and the time until the turbine recovers (Tr) increases substantially. If ΔPr is too small, the turbine will *never* recover and will stall. This is true for all variable speed wind turbines. At least some manufacturers (good data is scarce) have tended to make ΔPr rather large, in order to minimize both Area 2 and Tr. But there is a penalty on the grid for erring in either direction: recovery that is too fast (i.e. ΔPr too large) will starve the grid of power too soon, but recovery that is too slow (i.e. ΔPr too small) will increase the uncaptured energy and exacerbate the system frequency recovery. In the extreme case ΔPr too small will place the turbine at risk of stalling.

[36] IESO, Market Manual 2: Market Administration; Part 2.20: Performance Validation; Issue 5.0.

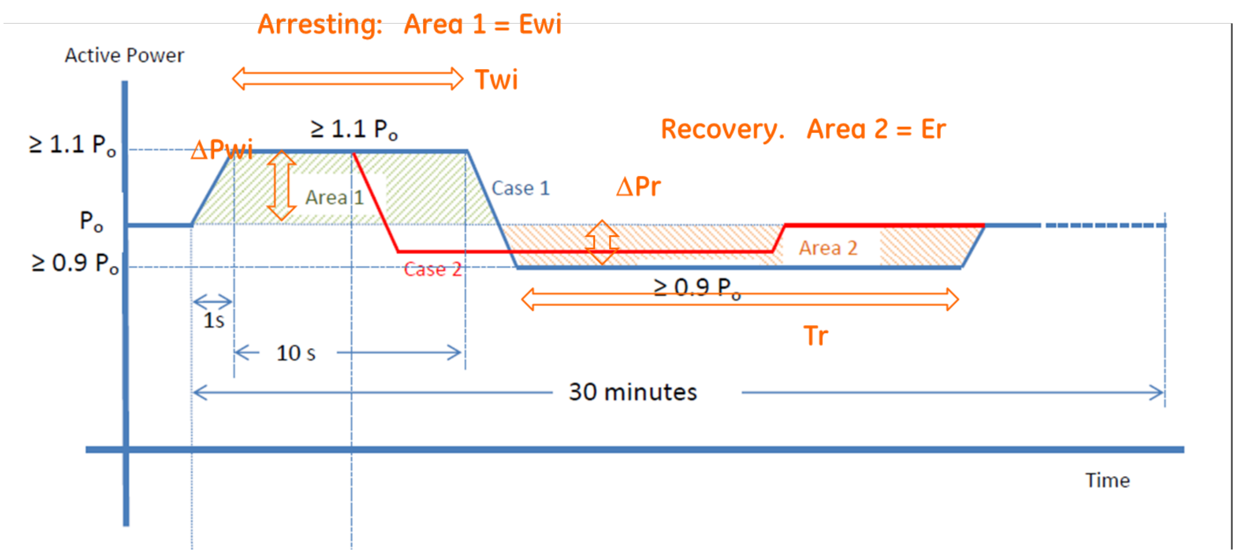

Figure 18 Open-loop FFR Control (based on draft IESO requirement)

The exact behavior will vary somewhat between manufacturers, and with the level of P_o. For calibration, the absolute minimum ΔPr that one OEM's turbines can tolerate is on the order of 7-8% of P_o.

The consequence of this complication is that, *in our opinion*, the inertia-based FFR should be designed to start "withdrawing" at the frequency nadir, and not wait until the frequency has actually recovered. Incurring the extra rotational energy deficit associated with continued inertia-based FFR after the nadir does not produce a significant performance benefit, as long as the release is not too abrupt. The open-loop control philosophy drives a very "bimodal" response from these resources: with no response at all for events below the trigger threshold (e.g. inside the deadband), and the same aggressive response for all events that meet the triggering criteria. Such abrupt discontinuities in power system control have been known cause problems.

The FFR control implemented by Enercon for Hydro Quebec[37] is a triggered response, which is released with some hysteresis, i.e. when the frequency recovers to a given level (0.05Hz) above the trigger frequency, it begins to ramp back over a period of time (5 seconds). The triggered release avoids unnecessarily sustained response. The results shown in Figure 19 are illustrative (there are several more good traces in the paper).

[37] Markus Fischer, Sonke Engelken, Nikolay Mihov, Angelo Mendonca, "Operational Experiences with Inertial Response Provided by Type-4 Wind Turbines", Wind Integration Workshop 13, 2014.

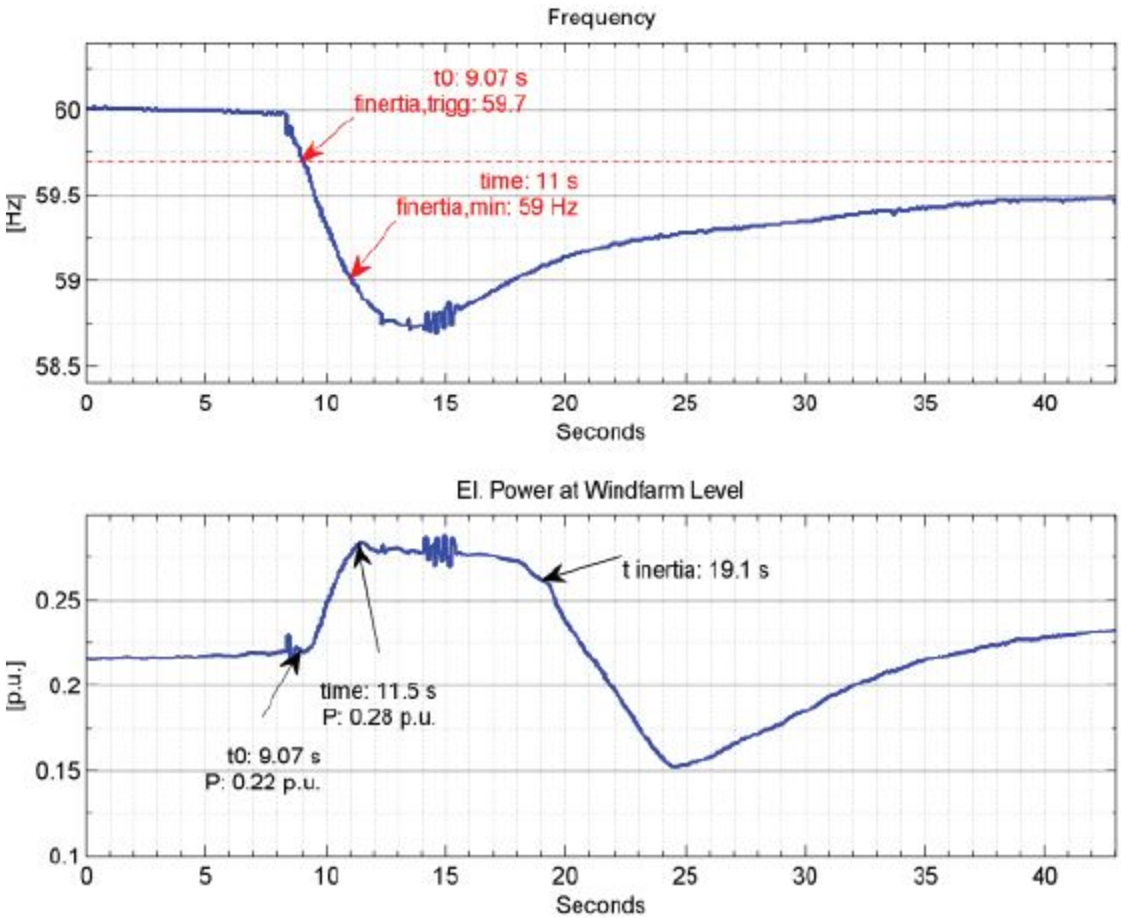

Figure 19 Field Measurements of inertia-based FFR from Quebec[37]

The closed-loop control addresses most of these worries, with a gradient of responses related to the error (i.e. to the amount that event passes the threshold for triggering). This allows the triggering threshold to be set closer to "normal" with less fear of adverse consequences from unintended triggering. But as noted above, very high closed-loop gains may raise stability worries. It is possible that a hybrid of open and closed loop controls is appropriate. For example, an open-loop "kicker" on the closed-loop control can discretely "force" the control to act aggressively in response to specific stimulus (e.g. a trigger from an SPS). This maintains the benefits of a closed-loop control, while making better use of the equipment capability for extreme events – like separation of South Australia from the NEM. There is precedence with HVDC, and in the downward direct for wind plants. Upward kick is NOT the present state of the art, although the industry in some places is evolving in that direction – i.e. very aggressive open loop response when frequency thresholds are passed.

3.2.3.3 Frequency Response via Torque Control on Type 2 Wind Turbines

Type 2 wind turbines are no longer being built by Vestas, but there are installations on the NEM. Like Type 1 machines, these machines are electrically coupled to the grid and are sensitive to system frequency. Standard induction generator equations apply. Figure 20 shows the behavior of a type 1 WTG in a system with rapidly dropping frequency and low inertia. The initial upward spike in power is the inertial response, and it is beneficial to the system. The recovery swing is not, in general, helpful.

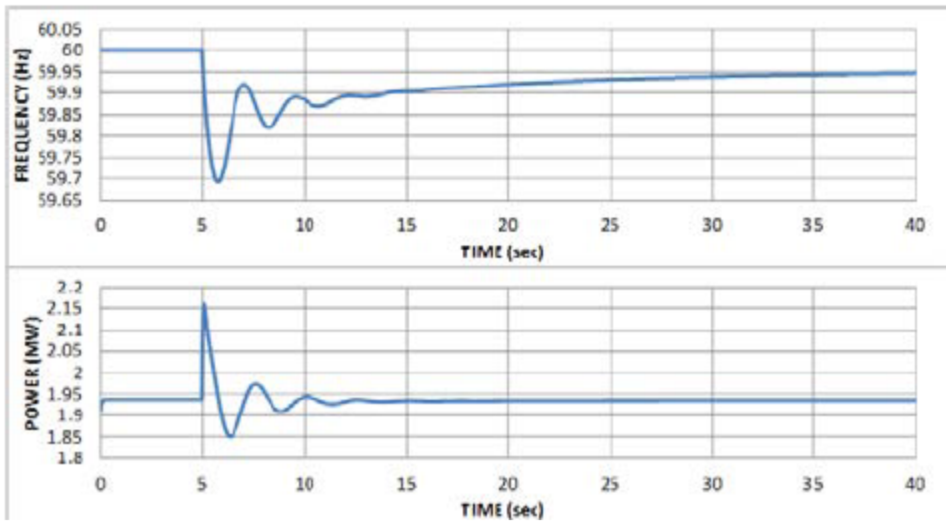

Figure 20 Frequency response for a Type 1 WTG in a system with low inertia[38]

In the case of a Type 2 machine, there is a degree of control gained by having fast control (via a power electronic chopper circuit) over the external field resistance. In Figure 21, a family of speed-torque curves is shown. The black trace represents the "normal" condition, with the external resistance (total rotor resistance is shown in the legend) shorted (only internal rotor resistance is left). The machine will see an increase in slip from the pre-disturbance level (suggested by the blue dot) as the grid frequency falls, increasing torque (as suggested by the blue arrow). Increased toque will cause more power to be delivered to the grid, a benefit similar to that from synchronous inertia. If the initial operating point were on one of the colored curves, the torque could be further increased, added to the arresting power provided by the WTG. This is essentially an FFR service, but it comes at the expense of having more rotor resistance, therefore more loses, and higher rotor speed. Higher rotor speed means faster blade tip speed, which may (or may not) reduce the power being captured from the wind. Therefore, there is the possibility of a non-trivial efficiency penalty associated with providing this control. Specifics of the turbine and generator would need to be considered.

There is also the possibility of using the field resistor to ease the power recovery after the frequency nadir. Once the WTG has slowed, as the grid frequency recovers post nadir, the

[38] Muljadi, E, et al – NREL "Understanding Inertial and Frequency Response of Wind Power Plants", July 2012 NREL/CP-5500-55335.

torque will drop allowing the rotor to recover the energy lost to the grid. This has similarities to the recovery of the type 3 and type 4 WTGs. The field resistor control will allow some control of the motor toque during this phase (as suggested by the red arrow), such that the recovery power (see the backswing in Figure 20) is less acute and the time of the recovery is extended. Like the other machines, there are aero-mechanical constraints that must be respected. It is not known whether this type of control has been commercially offered for type 2 WTGs.

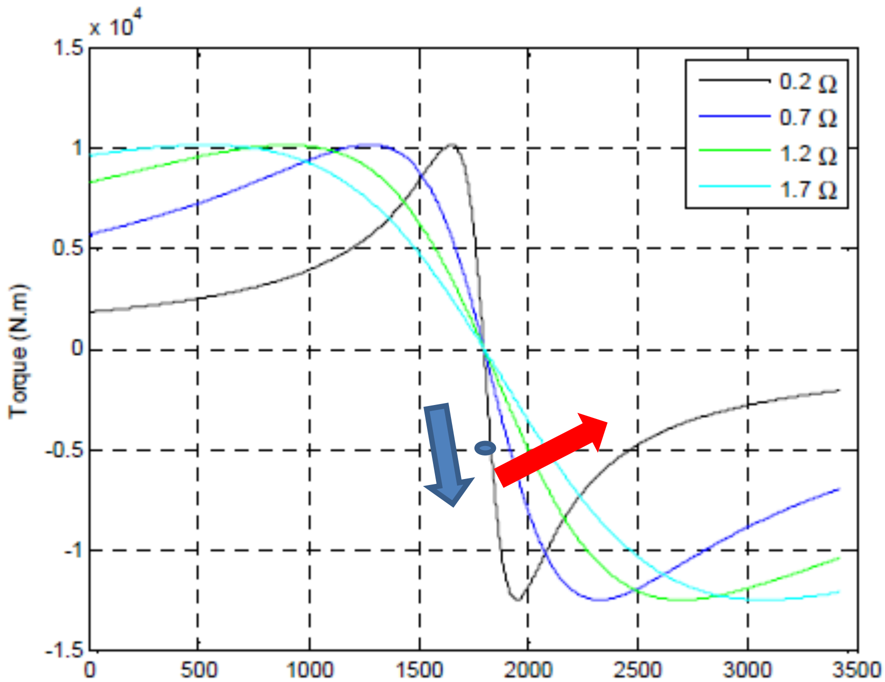

Rotor Speed (RPM)

Figure 21 Speed - Torque Curve NREL[39]

3.2.4 How Inertia-based Response Coordinates with other Wind Turbine Controls

Unlike the inherent response of synchronous machines, inertia-based WTG response is dependent on active controls. The primary control actuators are generator torque and pitch control. These actuators are also dedicated to meeting other control objectives for the WTG. Variations in wind speed, and particularly turbulence, constantly change the torque and other mechanical forces on the drive train, blades and entire structure of the turbine. The same actuators that provide inertia-based response are also relied upon to manage the turbine speed and mechanical stresses. These stress management controls take priority over inertial control. As noted above, rating of the turbine constituents must be respected. During more violent events, such as transmission system faults, functionality that provides fault ride-through capability to wind turbines will take precedence of FFR function. Measuring frequency

[39] Bastos, A.F., et al "Use of Newton's Method for Rotor-Resistance Control of Wind Turbine Generators" ICREPQ '12, Spain.

(much less responding to it with a specific active power response) during a severe fault is technically untenable for all technologies.

Further, the transient dynamic range of the actuators is dependent on the operating condition of the wind turbine. For example, at lower wind speeds the generator and converter will tend to have additional 'room' to increase power, BUT the speed of the turbine may be lower, meaning that there is less inertial energy stored and less speed margin to aerodynamic stall. Conversely, at wind speeds above rated, there is additional power available from the wind, so that extra energy (power) can be captured from the wind, rather than slowing the turbine drivetrain. However, the time dependent power ratings of all the components, especially the converter, must be respected, and these considerations apply to all OEMs, regardless of their approach to implementing Inertia-based FFR.

3.2.4.1 How WTG inertial response coordinates with pitch enabled Primary Frequency Controls

The active power controls offered by some wind OEMs typically work through the wind power plant's supervisory control system. The controls typically impose curtailment orders, ramp rate limitations, and PFR[40]. PFR is the autonomous frequency sensitive controls closely akin to governor controls for thermal and hydro generation. These controls also respond to significant deviations in grid frequency at the wind plant, increasing ("raise") or decreasing ("lower") power output in response to low or high grid frequency events, respectively. Slower controls that respond to external signals (usually from the grid operator), like Regulation or contingency dispatch, also work through the plant supervisory system. Unlike inertia-based FFR controls, these controls alter the active power control reference targeted by the turbine controls. Therefore, they can, under some circumstances, have significant impact on energy production. (Further discussion of primary frequency control is provided below in Section 4.4.) In the NEM, there are a range of FCAS's. These services are related to PFR, but have their own characteristics. In so far as these services provide active power response in proportion to frequency error, they are equivalent to PFR. While FCAS functions can clearly be provided by this type of control, it is not a given that IBFFR will seamlessly integrate with them. This is part of the necessary detailed analysis recommended.

When enabled, the response of the active power control to significant under-frequency events is to increase WTG output. Usually, the command for this response emanates from the wind plant level control, and is delivered to each individual WTG. To increase active power, the plant must be partially curtailed, so that additional power can be extracted from the available wind. This incremental power order signal will *add to* the power from inertia-based FFR control which is local to the individual WTG. The total response of the WTG to these two signals is coordinated to respect the physical capabilities of the WTG. Again, whether these controls reside in the plant level controller or in the individual turbines, depends on the manufacturer.

40 Cardinal, M.E; N.W. Miller, "Grid Friendly Wind Plant Controls: WindCONTROL – Field Test Results"; proceedings WindPower 2006, Pittsburgh, PA .

And, as noted above, approaches to these controls vary considerably by OEM, and not all vendors even offer this class of control.

It is important to reinforce that, in general, these active power controls reside at the plant supervisory level. This means that signals from the host utility (e.g. curtailment orders) as well as systemic measurements are processed at plant controller, and individual signals (usually power set points) are sent to the individual turbines in a wind plant. This has many performance and security benefits (e.g. the plant controller understands when individual turbines are unavailable, and can compensate), but it requires communication from the plant controller to the turbines. This takes time. Communications latencies vary with plant size (count of turbines), OEM and vintage. A reasonable range of communication latencies for present technology commercial wind plants is on the order of 200 to 500ms.

Grid over-frequency events present a different risk to system operations. Excessive high frequency is stressful to power components. Further, temporary high frequency swings can present a reliability concern. For example, in one recent well publicized grid event[41], the high frequency backswing from a major grid disturbance caused power plant trips and aggravated an already severe event. When enabled, the primary response (at the plant level control) will reduce power output for the duration of the over-frequency event. This behavior is similar to that of governor control on thermal generation, except that it is faster and allows deeper runback of power than is typical of thermal generation. Having this function has been mandated in EirGrid and ERCOT (Texas) for years. A broader discussion of backswing over-frequency is provided below, in Section 6.6.

3.2.5 Grid Code for Inertial Response for Wind Turbines

Ultimately, system operators have begun modifying grid codes to include some type of inertial response requirement. The development and demonstrations by several OEMs, shows that such functionality is, indeed, possible. However, it also shows that inertial response identical to that of synchronous generation is neither possible nor necessary. Controlled inertial response of wind plants is in some ways better than the inherent inertial response of conventional generators. Inertial response of wind generation is limited to large under-frequency events that represent security and continuity-of-service risks to the grid. The crafting of new grid codes should therefore proceed cautiously and focus on functional, systemic needs.

The following lists some elements that should be included (from our perspective) in sensible FFR technical requirements for wind turbines:

- Specification of minimum turbine power below which inertial response is not required
- Specification of preferred interactions/priority with other wind plant controls (E.g., should FFR or curtailment have priority? Should reactive support or frequency support have priority?)

[41] FRCC Event Analysis Team (FEAT) Interim Recommendations Report; http://www.balch.com/files/upload/FRCC_Interim_Report_6_3_08.pdf

- Not specified to be identical to synchronous machine
- Not specified to be identical for all operating conditions
- Not specified to be exactly reproducible with individual turbine tests
- Not specified to be energy neutral (for all events)
- Not specified to be overly prescriptive or requiring of impossible power recovery constraints
- Not specified to require delivery of energy beyond that necessary to improve the frequency nadir

3.2.6 Key Technology Considerations from the Perspective of WTG Manufacturers

This section is intended to provide some insight into how wind turbine manufacturers (OEMs) will view new performance requirements. It is, of a necessity, somewhat speculative and based on our experience.

3.2.6.1 Limits to Performance

What is it that OEMs will worry about when asked for very fast and aggressive inertia-based response? They will concentrate significantly on turbine speed and mechanical loading. In general pitch response isn't the limit; under the conditions when the turbines are pressed to meet performance objectives, the blades tend to be at pitch limits anyway. Even for combined performance (with governor/PFR behavior), the blades will tend to catch up. Likewise, the inverter itself is very fast. The limitation on response is much more related to interaction with other controls in the turbines. For example, very aggressive response stimulates the drive train dynamics. Active loads management in (all/most) modern WTGs will tend to respond to abrupt changes in torque, and these controls are likely to want to defeat the inertial-based FFR controls. Likewise, turbulence and loads management are considerations that take priority over grid functions. Therefore, demands that a certain level of response be provided under ALL conditions will inevitably be met with conservative response from OEMs.

This raises several interesting questions about requiring, providing and paying for performance. Basic reality is that if resources are rewarded for doing the best they can at any instant, the result will be better. It presents an interesting riddle: is this ("do the best you can, and you will be rewarded") an appropriate response to "non-credible" events (like loss of both circuits of the Heywood interconnector)? Is this a legitimate response to design-basis (credible) events? Or should AEMO handle the two classes of events differently? This is part of a bigger industry debate.

OEMs will worry a great deal about the frequency with which the turbines are subjected to stressful control actions. They will ask: How many times/year would the control trigger? For example, an OEM is likely to get nervous at about 100/year. Repeated actuation starts to have fatigue costs; these are very difficult to quantify, but will affect willingness to have aggressive response.

3.2.6.2 New Equipment vs Retrofits

The fact that these (FFR) controls will interact with a wide variety of incumbent controls makes retrofitting challenging. The reality is that each OEM will need to evaluate what is possible for each platform (and maybe even each model). Much like the HVDC discussion below, only the OEMs will really know what is possible. It is by no means guaranteed that they will know what is possible without making specific (and expensive) investigation within house. Don't assume that the OEMs "just know" this stuff, they probably don't. That has serious implications for either retroactive requirements or new A/S inducements: the owners of existing wind plants will not change their equipment, unless the incentive (carrot or stick) is big. High penalties for failure to respond to retroactive requirements will result in howls of protest.

3.2.6.3 Costs

The capital cost to wind plant owners for inclusion of inertia-based FFR controls (and PFR) capability in a new plant is expected to be on the order of less than 1% of the capital cost of the overall project. While this is a small incremental cost, it is important to note that margins and returns on investment in electric power markets are measured in quite small fractions of the overall initial capital cost of a project. A less than 1% increase in the capital cost of project might reasonably be regarded as a substantial sum by the investors.

As noted above, retrofits of inertia-based FFR controls to existing plants are more expensive than inclusion in new plants. While capital costs for retrofitting the capability in wind plants vary depending on the vintage and type of technology, a rule of thumb is that the capital cost for retrofits is on the order of less than 2% to possibly much greater for some OEMs. Again, this might be a small incremental cost, but in the case of mandatory retrofits, this incremental cost has not been built into the economic evaluation, justification, and financing for a project, and so the economic consequence to the plant owner is more significant. It is noted that plant owners in ERCOT, when faced with a mandatory retroactive requirement (for PFR capability), were reportedly very unhappy.

3.2.7 Synopsis of Wind Turbine FFR

These are short answers to the questions from AEMO RFP; detailed discussions are provided in the preceding text.

How fast can it respond? Activation starts with ones of cycles of triggering at the turbine level, and power can rise on the order of 20-30%/second.

How much can it inject or absorb? Incremental power is generally limited to 10% of pre-disturbance power.

How long can it sustain this response? Below rated wind speed, but near rated power, the maximum available arresting energy is on the order of 1 MW-sec/MW; this maximum available arresting energy drops roughly linearly with power level.

Is there a particular ramp or shape to the response? Ramp rate is limited to 20-30%/second.

What determines each of these limits? Is there potential for it to be adapted? Mechanical stresses and aero-mechanical stability. Pushing further might be possible, but could also cause loss-of-life in turbines.

Where, when and how have these capabilities been demonstrated? Various OEMs, including, at least, GE, Senvion, Enercon and Siemens have demonstrated various inertia-based FFR controls. Tests have been run in the US and Canada, at least.

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. For designs are presently offered, marginal costs are small. There is nearly, but not quite, zero lost energy production. And, generally available designs are intended to not use up equipment life.

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? First generation offerings have been available for several years from selected OEMs. There is considerable interest now and new requirements are emerging rapidly. More options and more field data will almost certainly be available in the near future.

How is the ability to deliver a fast frequency response service affected by nearby power system disturbances, such as transmission faults? Will the FFR service still be available in the period following a transmission fault? How could this be managed? Limitations are mostly systemic: ability to push extra power into the grid following disturbances is often voltage dependent. In weak grid controls, active power injection is often slowed or suppressed to allow for the grid to recover. This will be a serious concern for this application under conditions when loss of infeed is caused by a fault and trip event. It applies to all FFR technologies. A somewhat separate consideration is how fast the equipment can to resume power injection on AC voltage recovery. This is highly dependent on the state of the DC voltage. This is a design consideration for all FFR and will affect all the inverter-based technologies discussed in this report. , In general, technologies that resist changes in the DC bus voltage on the inverters (i.e. the DC bus is "stiff"), like batteries, will tend to be able to recover the fastest. (This reflected in the subsequent discussions and in the speed notes in the summary tables.)

What is the potential for false or spurious triggering? How is this managed? This is discussed in Section 1, and is common to all the available technologies.

Are there any other important considerations for the use of this technology to provide a fast frequency response service, to mitigate high RoCoF? See discussion.

3.3 Battery Energy Storage

Batteries convert chemical energy into electrical energy. An electrolyte allows ions to flow between the anode and cathode, while a DC electrical current feeds an external circuit. Batteries and hybrid battery systems (e.g., battery/supercapacitors) have the potential to provide a useful fast frequency response to mitigate high RoCoF. Some batteries have an extremely fast response time (e.g., lithium and advanced lead acid) and the technologies range from extremely mature (lead-acid) to less mature (flow). This report focuses on three promising technologies for FFR: lithium ion, flow, and advanced lead-acid.

In the last several years, batteries have been increasingly deployed in utility applications, through subsidies, mandates, and especially as demonstration projects for utilities and governments to better understand the capabilities and limitations of the technology. Applications include the provision of energy, capacity, ancillary services, transmission and distribution upgrade deferral, demand charge reduction and backup power. Depending on the application, different characteristics of energy, power, duration, speed, etc. may be needed. Batteries can provide value to participants across the power system spectrum. For example, storage can support the grid to provide capacity to meet system peak, while demand charge management may allow an individual customer to reduce his individual demand peak. Today, there are commercial opportunities for batteries in providing fast regulating (secondary) reserve in markets that pay for performance (e.g., PJM in the US). Batteries are also commercially viable in some applications and rate structures where it makes sense to help manage energy bills with existing demand charges and time-of-use rates. In 2013 SBC reported about 750 MW of large batteries deployed in the utility sector[42]. Today the US Department of Energy (DOE) Energy Storage Database lists over 731 battery energy storage systems (BESS) projects (including Lithium, lead-acid, flow, and sodium technologies) with 1720 MW of total capacity world-wide[43].

BESSs are comprised of a power conversion system (PCS) and a battery. Battery cells are stacked to make up battery modules of the appropriate voltage. These modules may be controlled with a battery management system (BMS) to protect the battery from over-charging, over-discharging, and thermal damage. The PCS has a bidirectional inverter that converts DC output from one or more strings of batteries to AC output and vice versa. For example, in a 8 MW (32 MWH) BESS in Tehachapi in Southern California Edison, 56 battery cells (60 Wh, 3.7 V) make up each 3.2 kWh module (52V). Eighteen modules are in each rack and there are 151 racks in each string of the BESS. Charging/discharging voltages of the strings are at 760-1050 V_{DC}. Each 2 MW (8 MWH) string has 8 BMSs and one controller. There are two inverter lineups in each 4 MW PCS which then connects to a 12.47 kV_{AC} grid. The two PCSs are operated by a master controller which will send real and reactive power commands and may provide different modes of operation.

Speed of Response

Batteries are not limited by the chemical response time (mass transfer dynamics of ion migration and depletion), but by the inverter and controls response times. Additionally, the location of the fast controls, critical sensors, and actuators will have the largest impact on time because plant level control loops are slow compared to the PCS. The plant level response time consists of plant level sensing, communication and dispatch to individual inverters.

---

42 SBC Energy Institute, "Electricity Storage", September 2013.

43 http://www.energystorageexchange.org/ accessed July 17, 2016.

Depending on when the event occurs during the control cycle, battery engineers conservatively estimate that the response is between 150-400 ms. Because of this, an extremely fast frequency response is better implemented directly on the inverters. The PCS regulator control loop may have a tens of ms response time. For example, in Hawaii in the US, requirements for extremely fast response have led to frequency droop implemented directly on the inverters so that the system begins corrective action within 10 ms, with a full response within 50 ms. There are drawbacks of PCSs that are required to act fast: PCSs that must respond within 20ms are maintained in a hot standby condition that has parasitic losses (2% of rated power)[44].

In addition to the BESS response time, total response time for a BESS FFR would include RoCoF detection/communication time as discussed in section 2.2.2. Whether a direct event detection scheme or a local RoCoF detection is used, the fastest options for this detection/communication time appear to be on the order of two cycles. The BESS options discussed in detail below have activation and full response time on the order of cycles, making this a potential option for providing arresting energy within the 250 ms needed.

3.3.1 Lithium Batteries

Lithium-ion and lithium polymer-type batteries use lithiated metal oxide as the cathode and carbon as the anode, with lithium salts as the electrolyte. There are various chemistries including lithium iron phosphate, lithium manganese oxide, lithium nickel manganese cobalt oxide, lithium titanate, and lithium nickel cobalt aluminum.

They are widely used in consumer electronics and hybrid/electric vehicles. They are also increasingly used for utility applications, especially to provide secondary reserves (frequency regulation) in markets that pay for performance (e.g., PJM in the US) and they make up the bulk of BESS chemistries being deployed today. To give a sense of lithium battery deployment, the DOE Global Energy Storage database, that catalogues storage projects around the world, lists 430 operational projects totaling 1149 MW of lithium BESS installations worldwide.

Low internal impedance leads to high efficiencies (90% and greater). They provide high energy and power density. Lithium batteries have good cycle life (depends on depth of discharge, and while typically 1000-10,000 cycles, can be up to several 100,000s cycles), and beyond that, end of life is marked by degraded performance due to increased internal impedance or decreased capacity. They also have relatively low self-discharge. Disadvantages are cost and need for battery management systems. High battery temperatures negatively impact cycle life so thermal derating in the battery management system limits charge and discharge power limits based on battery temperature. For example, a lithium BESS may have a charge/discharge limit equivalent to the power rating of the BESS when cell temperatures are

[44] EPRI and the US Department of Energy, "EPRI-DOE Handbook of Energy Storage for Transmission and Distribution Applications", 1001834, Palo Alto, CA, December 2003.

within limits. At high temperatures, charge/discharge limits may decrease to half of the power rating. Expected lifetime is 5-15 years.

How fast can it respond? Lithium BESS response times ($T_{Activate}$+ $T_{ActivateFully}$) are reported to be about 10-20 ms[45,46,47]. Documented response times (see below) for tests of lithium BESSs in the field show tens of ms to hundreds of ms, but we note that these systems were not designed for the extremely fast responses needed in South Australia. Discussions with industry on projects that may push the envelope on speed indicate that total response time ($T_{TotalResponse}$) can be done in a few cycles. The very fast responses reported for batteries in this are, in part, likely due to the stiffness of the DC bus, as discussed above.

How much can it inject or absorb? Lithium batteries can inject or absorb 100% of their capacity but this impacts cycle life. A123 (now NEC) reported a 100,000 cycle lifetime at 100% depth-of-discharge (DOD)[48].

How long can it sustain this response? Response speed is not a driver for thermal stress, making lithium batteries well-suited for FFR. However, for a system designed for high power, the magnitude and duration of response over the time frame of hours may lead to constraints due to internal resistance and mass transfer dynamics that may limit longer, sustained response. Limitations on response and proper BESS design and cooling capacity manage various applications' needs. But for the timeframe of FFR, which may require sustained response for up to 15 minutes, lithium BESSs are suitable.

Is there a particular ramp or shape to the response? No.

What determines each of these limits? Is there potential for it to be adapted? The dominating factor of response time of the lithium BESS is the PCS. The battery cell response time is nearly instantaneous (tens of microseconds). There are tradeoffs in lifetime with DOD and sustain of response. Managing thermal impacts means there may be impacts in terms of efficiency for parasitic loads and in terms of capital cost with increased cooling capacity for higher performance.

Where, when and how have these capabilities been demonstrated? Over one hundred projects with total capacity of over 800 MW of lithium BESSs have been installed at the utility-scale for frequency regulation services[49]. While frequency regulation calls for responses on

---

[45] SBC Energy Institute, "Electricity Storage", September 2013.

[46] C. Vartanian, A!23, *Grid Stability Battery Systems for Renewable Energy Success*, www.neces.com

[47] HDR Engineering, Inc., "Update to Energy Storage Screening Study for Integrating Variable Energy Resources within the PacifiCorp System," Salt Lake City, UT, July 9, 2014.

[48] HDR Engineering, Inc., "Update to Energy Storage Screening Study for Integrating Variable Energy Resources within the PacifiCorp System," Salt Lake City, UT, July 9, 2014.

[49] http://www.energystorageexchange.org/ accessed July 17, 2016.

the order of seconds instead of tens to hundreds of milliseconds, it gives an indication of the level of deployment of lithium BESSs for fast response.

An example of a lithium-ion BESS providing fast response to stabilize frequency is on the Big Island of Hawaii in the US. A 1 MW/250 kWh BESS was commissioned at the Hawi wind plant on the Hawaii Electric Light Company (HELCO) grid in January 2013. On this 180 MW system, variability at two wind plants has led to strong frequency fluctuations. A GE study led to the installation of a BESS with funding from the US Department of Energy. Two real-time control algorithms were developed: one to balance generation and load within 100 ms and another to smooth strong variations in wind within 200 ms[50]. Frequency variability has been reduced by 30-40% with the BESS. The frequency regulation algorithm sets BESS output as a function of frequency (not RoCoF). Response times were fast, comprised of the site dispatch controller's nominal computation time of 100 ms and 100 ms for the BESS's output to reach 90% of the command. A test of this latter capability is shown in Figure 22.

Another example, from the first Chilean BESS used to provide primary frequency response, was installed in 2009. This 12 MW/4 MWh lithium nanophosphate battery system controls output as a function of frequency (not RoCoF): outside of a +/- 0.3 Hz deadband it responds to overfrequency with a full lower, and to underfrequency with a full raise. While full power response times have been as fast as tens of ms, typical latencies are in the hundreds of ms[51].

Another example is the 500 kWh BYD lithium-iron phosphate battery in Qatar, with a < 100 ms response time. This system provides reactive power support, frequency regulation and black start capabilities[52,53].

---

50 Hawaii Natural Energy Institute, *Development of Real-time Closed –loop Control Algorithms for Grid-scale Battery Energy Storage* Systems, Aug. 2014.

51 Hsieh, E. and Johnson, R. "Frequency Response from Autonomous Battery Energy Storage", CIGRE US National Committee 2012 Grid of the Future Symposium, Paris, 2012.

52 IRENA, "Case studies: battery storage", http://www.irena.org/DocumentDownloads/Publications/IRENA_Battery_Storage_case_studies_2015.pdf

53 BYD, http://www.byd.com/energy/ess.html.

Batteries and National Grid's Enhanced Frequency Response

The U.K.'s National Grid has established an Enhanced Frequency Response (EFR) service. A tender for 200 MW of EFR was issued in July 2016. EFR acts when frequency is outside of a 49.5 – 50.5 Hz deadband. The time for detection, signaling, and delivering the MW output must be no greater than 1 second in total, with the time for detection and signaling to be no greater than 500 ms. Full output should be sustained for at least 15 minutes. While the auction is technology-neutral, battery developers are highly interested in providing EFR, as demonstrated by the 888 MW of battery projects (out of a total of 1.3 GW of projects) that responded to a call for expressions of interest. Renewable Energy Systems (RES) is already working with National Grid outside of this auction to deploy 20 MW of battery storage that can provide sub-second frequency response. This is expected to be operational by the end of 2017.

Firm Frequency Response is divided into low frequency primary (full output within 10 s and sustained for another 20 s), low frequency secondary (full output within 30 s and sustained for 30 min), and high frequency (full output within 10 s and sustained indefinitely). The largest installed battery system in the UK is the Smarter Network Storage project 6 MW (10 MWh) which uses Samsung SDI batteries and S&C Electric inverter systems. This system relieves thermal overhead line constraints and also transformer constraints at the Leighton Buzzard distribution substation. It provides peak shaving, frequency response, reserves, and arbitrage. The graphs below show (top) testing of the system providing dynamic firm frequency response and (bottom) the system providing static firm frequency response.

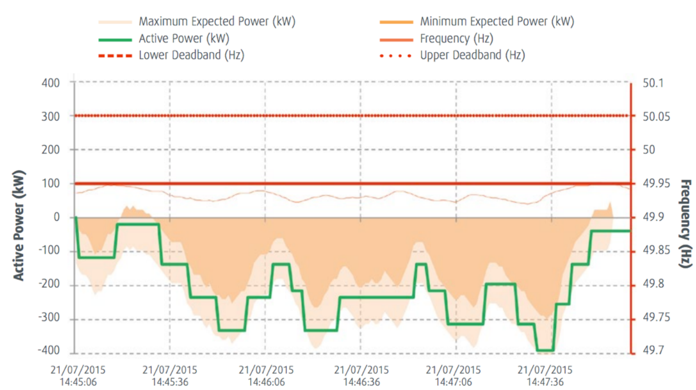

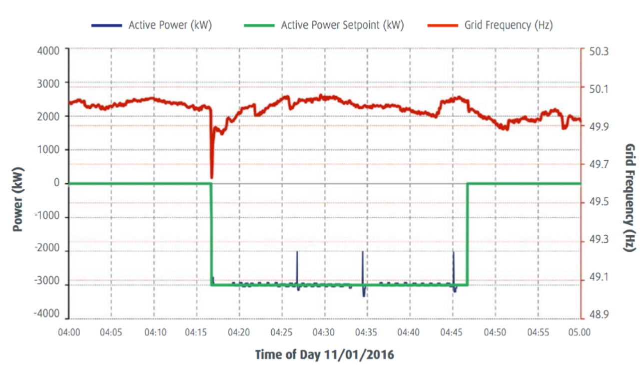

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. For a lithium BESS to provide FFR with a sustain of 15 minutes or less, there is no impact on the resource lifetime or increase in short run marginal cost. That being said, cycle life (typically 1,000-10,000 cycles) is an issue for lithium BESS. Capacity is typically degraded in 5-10 years and while those batteries may be used in other applications, they may need to be replaced to continue to provide FFR.

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? Lithium batteries are not as mature as lead-acid batteries, but deployment is rapid. Their development has been driven by consumer electronics and they are the technology of choice for electric vehicles. Of the BESS technologies currently being deployed in utility applications, lithium dominates the market. Their relative immaturity combined with wide deployment means that costs are dropping quickly and have a great potential to decrease further. With increased interest in BESS for FFR in Ireland, ERCOT, the NEM and other jurisdictions, it is likely that speed and lifetime can improve as well in the near future.

What is the potential for false or spurious triggering? How is this managed? As with the other inverter-based technologies in this section, this depends on the RoCoF detection scheme (see section 2.2.2.2).

Are there any other important considerations for the use of this technology to provide a fast frequency response service, to mitigate high RoCoF? Safety has been an issue with lithium-ion batteries, with overheating being a potential problem. $LiFePO_4$ is one of the two commercial lithium chemistries that is not subject to thermal runaway and therefore considered by some to be safer.

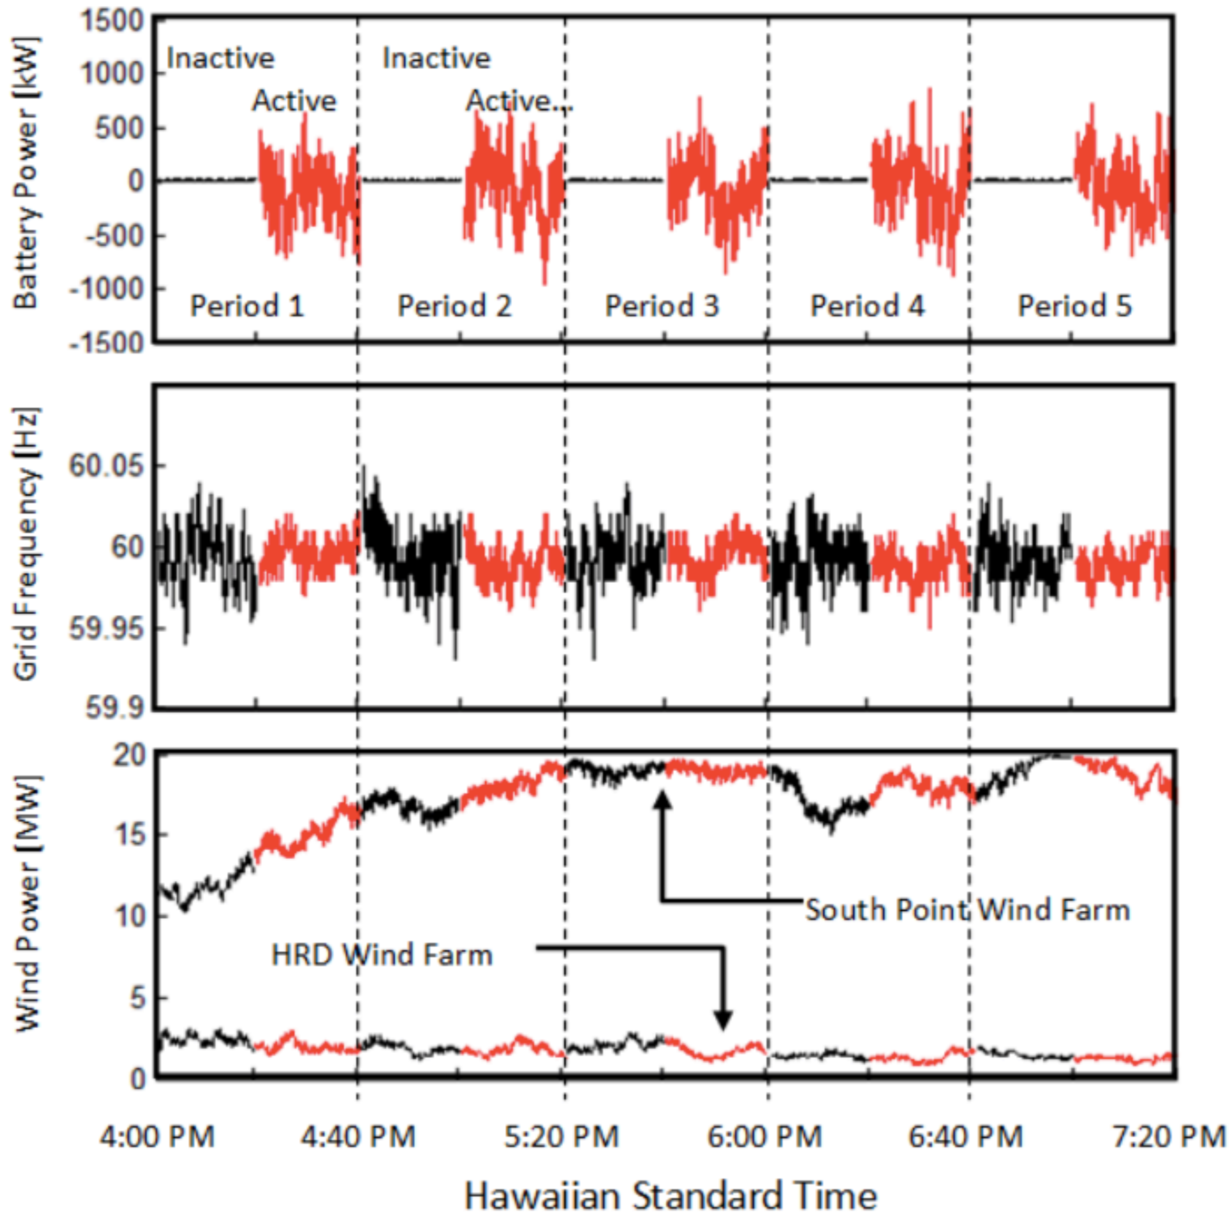

Figure 22 Performance of HELCO BESS using frequency regulation algorithm[54].
The black segments show battery power, grid frequency and wind power when the BESS is inactive; the red segments show the same parameters when the BESS is active.

3.3.2 Flow Batteries

Flow batteries are similar to fuel cells. They have active materials in solution in the electrolyte at all times. Two separate electrolytes are separated by a membrane which allows ion

[54] Hawaii Natural Energy Institute. *Development of Real-time Closed-loop Control Algorithms for Grid-scale Battery Energy Storage Systems*, Aug. 2014.

permeability, while current flows through an external circuit. The electrolyte is pumped past the membrane and external tanks hold some amount of electrolyte, which determines the energy capacity rating of the battery. A commercial flow battery technology is the vanadium reduction and oxidation (redox) battery (VRB) which is based on different ionic states of vanadium. A proton exchange membrane allows ions to flow and a vanadium/sulfuric acid mixture is used as the electrolyte for both sides of the cell. Other chemistries include zinc bromine, hydrogen bromine and zinc nickel oxide.

The DOE Global Energy Storage database lists 67 operational flow BESS projects worldwide, totaling 74 MW[55]. An advantage of flow batteries is that power and energy can be scaled independently: power ratings depend on the surface area of the electrodes while energy capacity depends on the volume of electrolyte in the tanks. This makes flow batteries attractive for long duration storage. Flow batteries also have the ability to use 100% of the stored energy. They are tolerant of over-charging, over-discharging and wide temperature ranges. The electrolytes are stored in separate tanks so self-discharge is minimal[56].

Another advantage is the long cycle life and lack of degradation of power or energy ratings over the lifetime of the battery. This is because cycling a flow battery does not cause a physical change in the electrode, as it does in most other battery technologies. Cycle life is thus independent of depth of discharge, in contrast to other batteries. Lifetime is limited by the cell stack and other components (tank, power electronics)[57].

Disadvantages include the more complex nature of the system. Flow batteries do not have high efficiencies. VRB efficiencies are on the order of 60-85%[58].

How fast can it respond? If the VRB is in a 'ready' state, with its stacks primed with reactants, they can increase from zero to full output ($T_{Activate}$+ $T_{ActivateFully}$) within a few milliseconds[59]. Short duration discharges can be responded to without even running the pumps[60]. Response time is thus limited by the controls and communications equipment, similar to other batteries.

For example, Prudent Energy's VRB system claims 100,000 cycles lifetime at full rated depth of discharge and a step response from charge to discharge of less than 50 ms[61]. UniEnergy Technologies' VRB system claims unlimited cycles over the 20 year life of the system, with

---

55 http://www.energystorageexchange.org/ accessed July 17, 2016.

56 UET specifies less than 2% self-discharge due to residual electrolyte left in stacks, with no discharge of energy remaining in electrolyte tanks over time. https://www.nwppa.org/wp-content/uploads/UET-The-Only-Megawatt-Scale-Containerized-Flow-Battery-NWPPA-presentation-05-17-16.pdf

57 Akhil, A., Huff, G., Currier, A., Kaun, B., Rastler, D., Chen, S. Cotter, A., Bradshaw, D., Gauntlett, W., "DOE/EPRI 2013 Electricity Storage Handbook in Collaboration with NRECA", SAND2013-5131, July 2013.

58 SBC Energy Institute, "Electricity Storage", September 2013.

59 Akhil, A., Huff, G., Currier, A., Kaun, B., Rastler, D., Chen, S. Cotter, A., Bradshaw, D., Gauntlett, W., "DOE/EPRI 2013 Electricity Storage Handbook in Collaboration with NRECA", SAND2013-5131, July 2013.

60 ibid.

61 Prudent Energy 2011 brochure.

100% available state-of-charge and a response time ($T_{Activate}$+ $T_{ActivateFully}$) of less than 100 ms[62].

How much can it inject or absorb? Flow batteries can easily inject or absorb 100% of their capacity. Additionally the cell stack can produce three times rated power output if the state of charge is between 50 and 80 percent[63].

How long can it sustain this response? An advantage of flow batteries is their ability to scale energy. Flow batteries can sustain their response for as long as they have capacity (which is a function of how much electrolyte is in the tanks) and there is no limitation on capacity.

Is there a particular ramp or shape to the response? No.

What determines each of these limits? Is there potential for it to be adapted? Similar to other BESSs, technical limitations on speed are determined by the PCS. The total response time is also determined by the RoCoF detection and communications time. The flow battery itself is not the limiting factor.

Where, when and how have these capabilities been demonstrated? Flow batteries have been deployed around the world, especially in Asia, and typically for demonstration or pilot purposes. There are far fewer flow batteries providing fast regulation or other fast services, compared to lithium. Of the 67 operational flow BESS projects in the DOE Global Energy Storage database, six projects totaling 5 MW were providing frequency regulation services[64]. In Japan, Hokkaido Electric Power Company installed Sumitomo Electric Industries' 15 MW/60MWh vanadium redox system to help manage wind and solar variability[65][66]. In Shenyang, China, Guodian Longyuan Windpower Company installed a Dalian Rongke Power's 5 MW/10 MWh vanadium redox battery system at the Faku wind plant to smooth wind output, regulate frequency and support voltage[67][68]. In Zhangbei, China, the China Electric Power Research Institute installed Prudent Energy's 500 kW/1MWh vanadium redox battery system at a wind and PV plant to provide fast balancing services[69][70]. In the US in April 2012, Prudent Energy installed a 0.6 MW behind-the-meter vanadium redox battery system in a California commercial process plant for demand charge reduction and to manage energy bills. In Pullman, WA in June 2015, Avista Utilities commissioned a 1 MW/3.2MWh vanadium redox

62 http://www.cesa.org/assets/Uploads/ESTAP-Slides-6.19.14.pdf

63 Mears, L., Gotschall, H., Key T., Kamath, H., "EPRI-DOE Handbook of Energy Storage for Transmission and Distribution Applications", EPRI, Palo Alto, CA, 2003.

64 http://www.energystorageexchange.org/ accessed July 17, 2016.

65 http://www.energystorageexchange.org/projects/2043 accessed July 17, 2016.

66 http://renewables.seenews.com/news/japan-s-hepco-sei-kick-off-15-mw-battery-system-verification-507909

67 http://www.energystorageexchange.org/projects/1624 accessed July 17, 2016.

68 http://www.rongkepower.com/index.php/article/show/id/140/language/en

69 http://www.pdenergy.com/pdfs/CEPRIProject-FactSheet053112-FINAL.pdf

70 http://www.energystorageexchange.org/projects/159 accessed July 17, 2016.

battery for load shifting, frequency regulation and conservation voltage regulation on a distribution feeder.

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. There is no impact on the resource lifetime or increase in short run marginal cost for the aforementioned level of performance.

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? Flow batteries are not as mature as lead-acid batteries and deployment is slower than for lithium batteries. While they have been around longer than lithium, they are used infrequently and don't really have applications in transportation or consumer electronics, which might decrease costs and increase deployment.

What is the potential for false or spurious triggering? How is this managed? As with the other inverter-based technologies in this section, this depends on the RoCoF detection scheme (see section 2.2.2.2).

Are there any other important considerations for the use of this technology to provide a fast frequency response service, to mitigate high RoCoF? Flow batteries have relatively low energy and power densities and therefore require more space than other BESSs. The balance of plant for flow batteries is also a bit more complicated.

3.3.3 Lead-acid Batteries

Lead-acid batteries are the oldest, most mature, commercial battery storage technology. They are used worldwide to start automobile and boat engines. Lead-acid batteries are the cheapest in terms of energy storage. While they can be discharged quickly, charging is slow, and aging leads to sulfation of the negative electrode, degrading performance.

Advanced lead-acid carbon batteries add carbon to the negative plate in various ways to increase power performance and handle partial states of charge. Some variants essentially add a supercapacitor to the lead-acid battery, such as the one shown in Figure 23, thus allowing the battery to provide an extremely fast response while also allowing for long duration time.

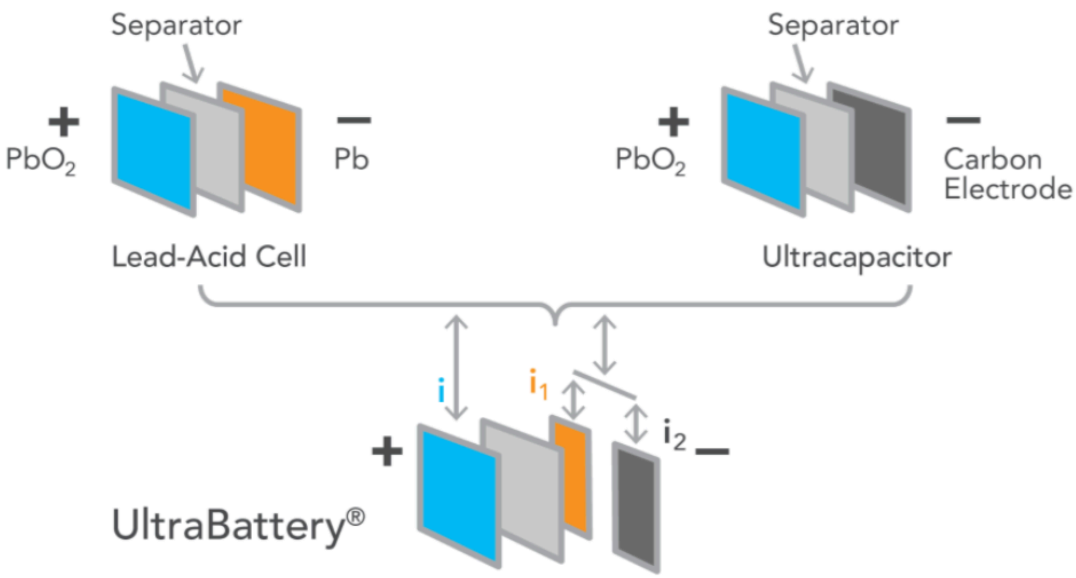

Figure 23 The Ecoult UltraBattery (hybrid supercapacitor and a lead-acid battery)

Advanced lead-acid carbon batteries can provide a high charge and discharge rate similar to that of lithium-ion batteries. Manufacturers report high cycle life (e.g., 1000-1600) with deep-discharges[71].

How fast can it respond? Figure 24 and Figure 25 show the fast response time, $T_{Activate}$+ $T_{ActivateFully}$, (<200 ms for full response) of a 5 MW (2.5 MWH) conventional lead-acid BESS in Vernon, CA that was installed in 1995 to provide backup power for critical processes at an industrial facility. Those tests are of considerable interest for AEMO, as they represent successful separation from the main grid (1st figure), and successful resynchronization with the main grid (2nd figure). The separated **system has no synchronous generation** of any kind running during this test. The BESS control here is responsible for *all* frequency control for the island. Consequently, while this demonstrates the capability of a BESS to respond, the example is not identically equal to FFR. Zero synchronous generation operation is discussed further in section 4.5.

[71] Akhil, A., Huff, G., Currier, A., Kaun, B., Rastler, D., Chen, S. Cotter, A., Bradshaw, D., Gauntlett, W., “DOE/EPRI 2013 Electricity Storage Handbook in Collaboration with NRECA”, SAND2013-5131, July 2013.

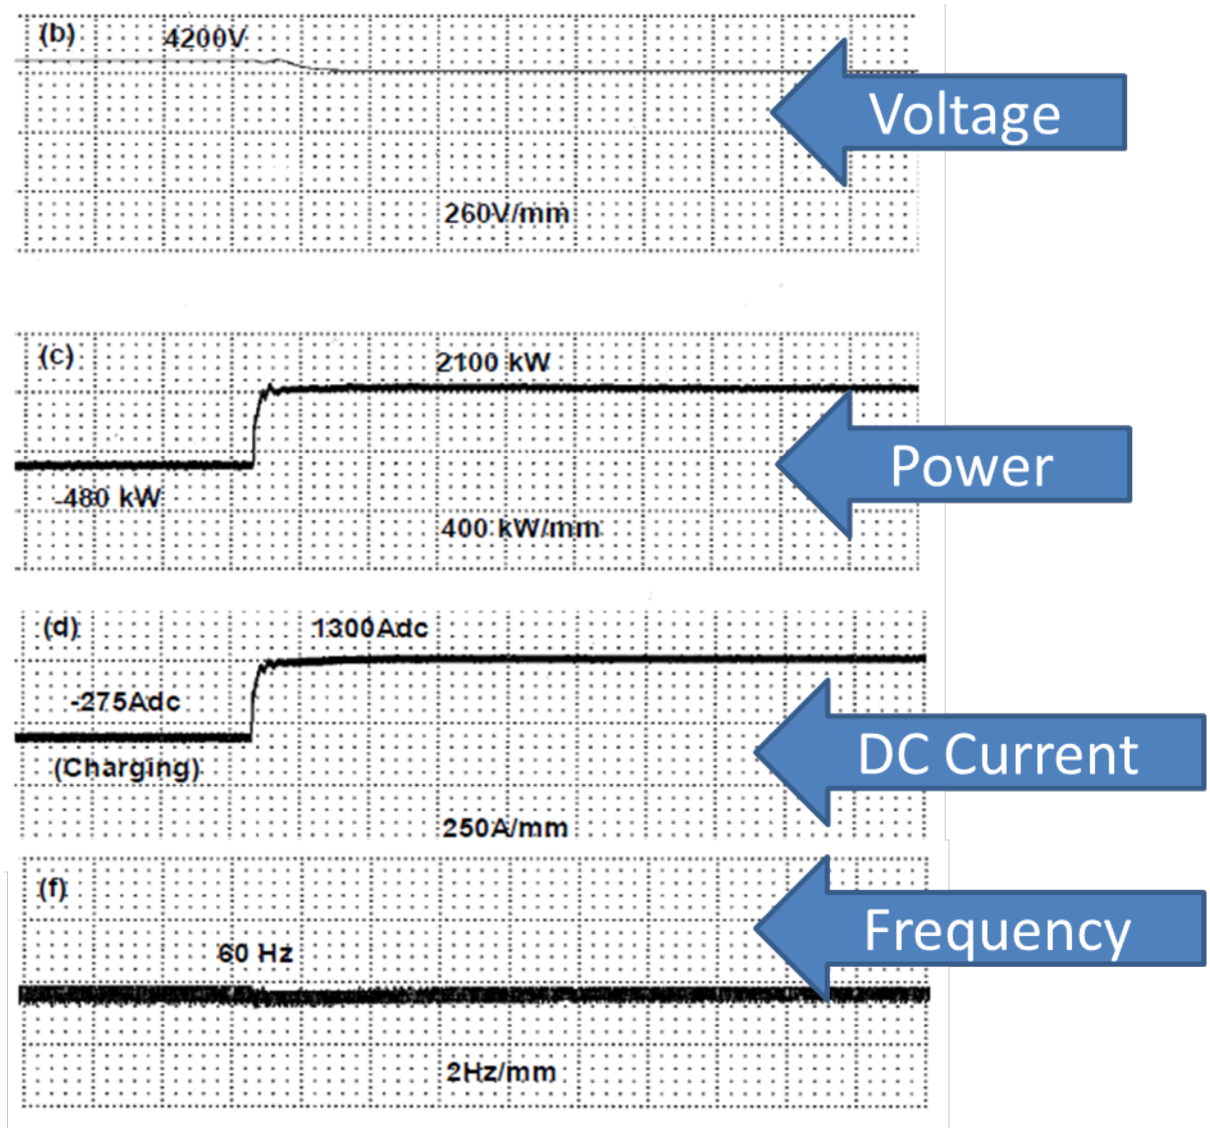

Figure 24 BESS supporting a trip to island event[72].

The x-axis is time, with each square representing 1 second.

[72] Miller, N. W., et al., "Design and Commission of a 5 MVA, 2.5 MWH Battery Energy Storage System, IEEE SPM 1996.

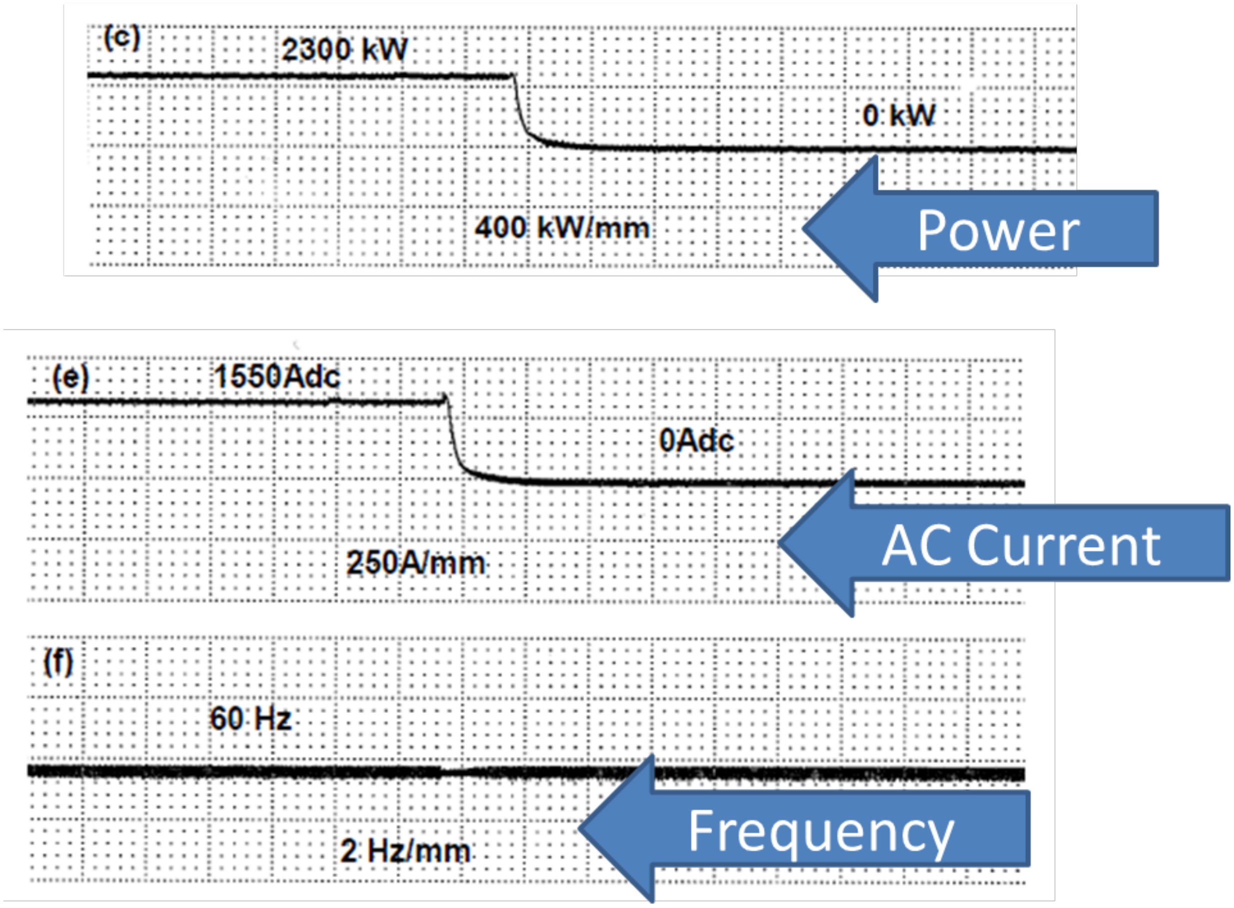

Figure 25 BESS enabling resynchronization of island to main grid[73].

The x-axis is time, with each square represens 1 second.

As with other BESS chemistries, total response times for advanced lead-acid BESSs are the sum of RoCoF detection time, communication time and response time of technology. Advanced lead-acid BESSs such as the UltraBattery (supercapacitor/battery) shown above are reported to respond in about 40 ms[74].

Figure 26 shows the performance of an advanced lead-acid BESS from Xtreme Power that was installed in ERCOT in 2013. Response time for this BESS to detect the issue (under-frequency in this application, not RoCoF), and fully respond ($T_{TotalResponse}$) is within 92 ms[75]. (This

[73] Miller, N. W., et al., "Design and Commission of a 5 MVA, 2.5 MWH Battery Energy Storage System, IEEE SPM 1996.

[74] HDR Engineering, Inc., "Update to Energy Storage Screening Study for Integrating Variable Energy Resources within the PacifiCorp System," Salt Lake City, UT, July 9, 2014.

[75] Duke Energy, "Technology Performance Report: Duke Energy Notrees Wind Storage Demonstration Project: 2015 Final Report", Nov. 2015. https://www.smartgrid.gov/files/OE0000195_Duke_FinalRep_2015.pdf

response is based on local measurements, and is subject to the concerns raised above in Section 2.2.2)

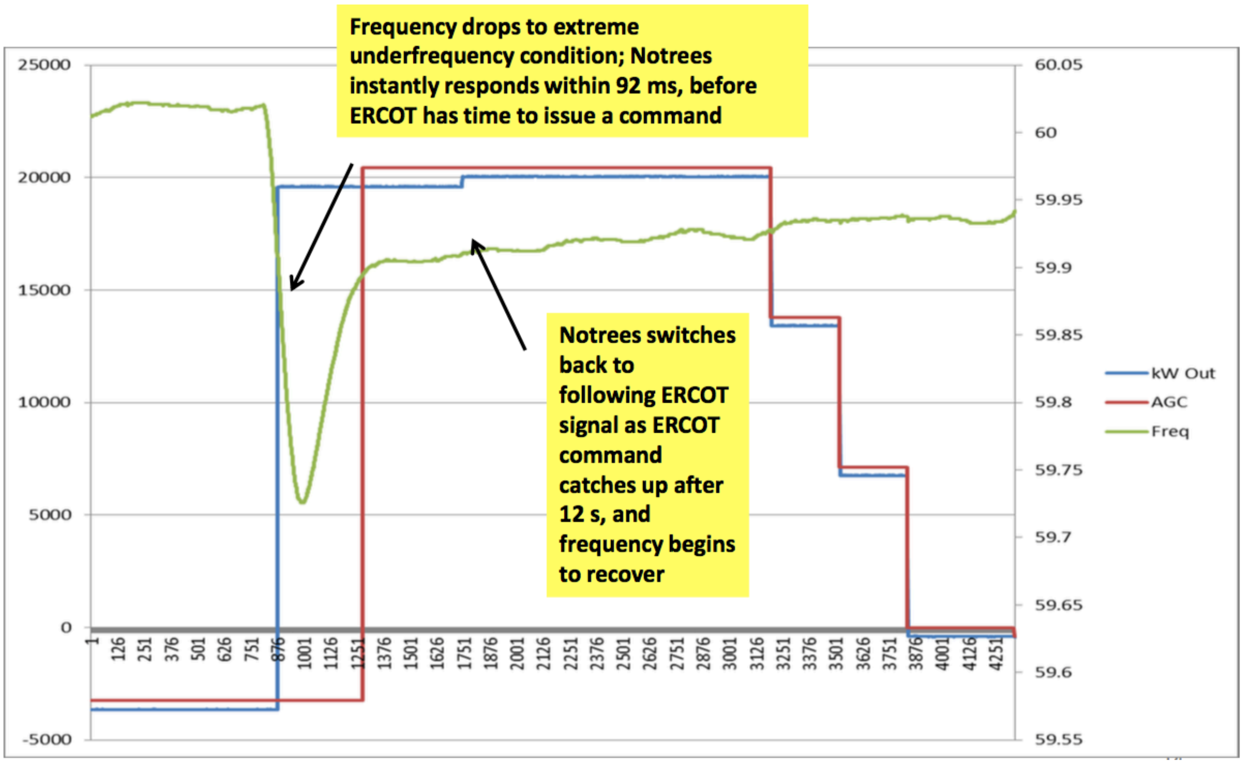

Figure 26 Notrees providing fast responsive reserves in ERCOT on Nov 1, 2013[76].

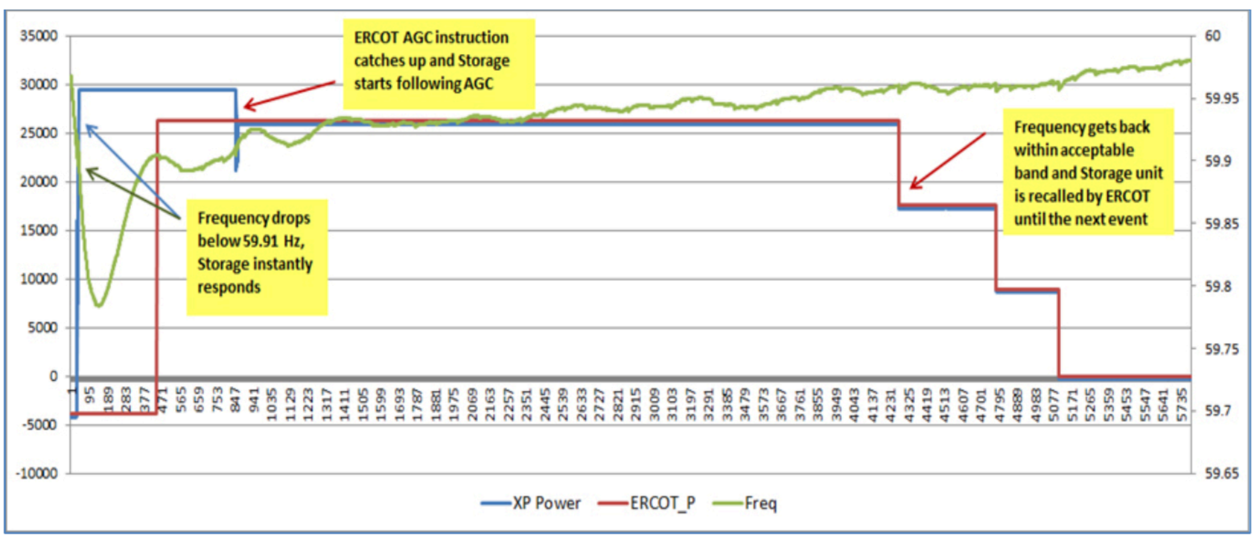

Figure 27 BESS arrests under-frequency deviation in ERCOT[77].

76 Ibid.

77 Ibid.

Left axis is energy in kW for the XP Power BESS (blue trace) and ERCOT AGC signal (red trace); right axis is frequency in Hz for ERCOT (green trace).

In Figure 27, the BESS responds nearly instantly to frequency dropping below 59.91 Hz. The BESS response gives the AGC time to ramp secondary reserves up.

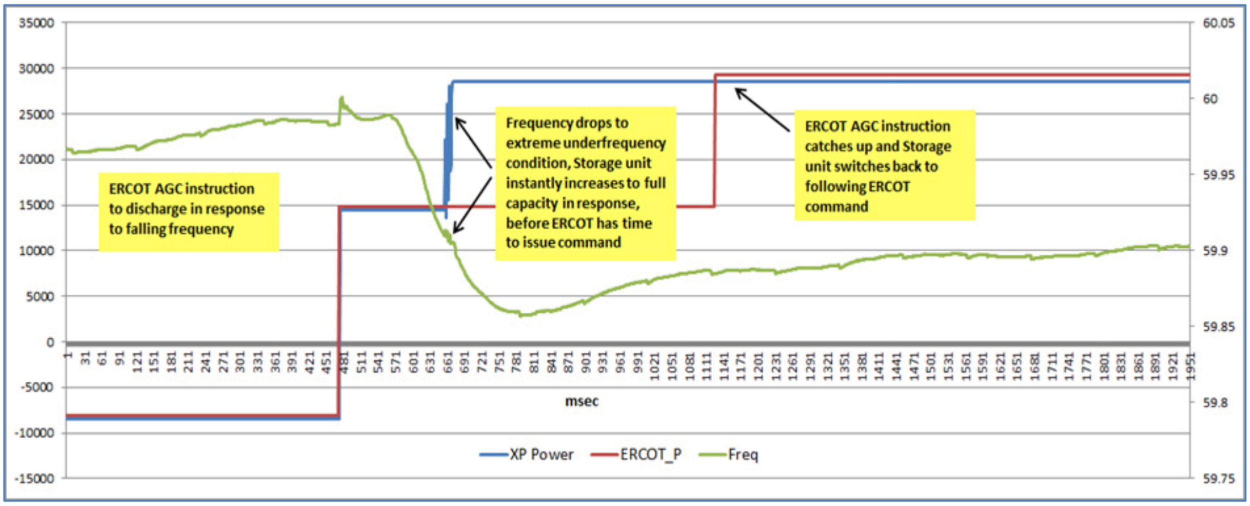

Figure 28 BESS arrests extreme frequency deviation in ERCOT[78].

Left axis is energy in kW for the XP Power BESS (blue trace) and ERCOT AGC signal (red trace); right axis is frequency in Hz for ERCOT (green trace).

Figure 28 shows another event with greater time resolution demonstrating the ability of the BESS to trigger and provide full response for under-frequency within cycles.

How much can it inject or absorb? Typical deep discharge lead-acid batteries may have a depth of discharge (DOD) of 60-80% but there are 100% DOD batteries developed for utility applications: Xtreme Power reported 1000 cycles at 100% DOD[79]. Ecoult's UltraBattery reported 9000 cycles at 100% DOD. Higher DOD negatively impacts cycle lifetime.

How long can it sustain this response? Lead-acid batteries can sustain a response for hours, depending on energy rating.

Is there a particular ramp or shape to the response? No.

What determines each of these limits? Is there potential for it to be adapted? The development of advanced lead-acid batteries that combine supercapacitors with batteries or alter the negative electrode give a fast and sustained response with 100% DOD. There are

---

[78] Ibid.

[79] HDR Engineering, Inc., "Update to Energy Storage Screening Study for Integrating Variable Energy Resources within the PacifiCorp System," Salt Lake City, UT, July 9, 2014.

also hybrid technologies that combine batteries with flywheels to give a similar high performance.

Where, when and how have these capabilities been demonstrated? The DOE Global Storage database lists 83 operational lead-acid BESS projects totaling 110 MW[80]. Fifteen of these are used for frequency regulation applications. Of these BESSes, 30 projects totaling 70 MW utilize advanced lead-acid batteries.

Before bankruptcy, Xtreme Power deployed several of their advanced lead-acid PowerCell systems. A number of these were installed in Hawaii to provide ramp rate control, voltage support, wind/solar smoothing, and/or reserves. For example, the 10 MW (20 MWH) BESS in Maui integrates a wind plant and the 1.1 MW (0.5 MWH) BESS in Lanai in 2011 integrates a PV plant. The Xtreme Power BESS shown in figures above was extensively tested and documented. This BESS was installed in the Notrees wind plant and provided fast frequency regulation in ERCOT[81]. The batteries, which had been commissioned Jan 2013, were no longer operating at full capacity by Oct 2015. Operations of that BESS had to shift to reflect the degraded performance and the batteries are being replaced with lithium ion in 2016.

A 3 MW East Penn UltraBattery (supercapacitor/battery) was commissioned in 2012 to provide frequency regulation services in PJM[82].

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. There is not an impact from the speed or the sustain of the performance. There is a tradeoff between DOD and cycle lifetime, as with most batteries (not flow batteries).

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? Lead-acid batteries are a very mature technology but their adaptation to advanced lead-acid batteries to provide fast services in the utility sector is relatively new, giving them potential for further improvement in the near term.

What is the potential for false or spurious triggering? How is this managed? As with the other inverter-based technologies in this section, this depends on the RoCoF detection scheme (see section 2.2.2.2).

Are there any other important considerations for the use of this technology to provide a fast frequency response service, to mitigate high RoCoF? Lead is toxic, but recycling is well-known and understood.

---

80 http://www.energystorageexchange.org/ accessed July 17, 2016.

81 Duke Energy, "Technology Performance Report: Duke Energy Notrees Wind Storage Demonstration Project: 2015 Final Report", Nov. 2015. https://www.smartgrid.gov/files/OE0000195_Duke_FinalRep_2015.pdf

82 Akhil, A., Huff, G., Currier, A., Kaun, B., Rastler, D., Chen, S. Cotter, A., Bradshaw, D., Gauntlett, W., "DOE/EPRI 2013 Electricity Storage Handbook in Collaboration with NRECA", SAND2013-5131, July 2013.

3.4 Flywheel Energy Storage

3.4.1 Flywheel Energy Storage Systems (FESS) with Inverter Interfaces

The relatively new generation of inverter interfaced flywheels are a commercially available technology with advantages of a high power density and high cycle life. Disadvantages include a low energy density.

These devices have two types of electrical topology that closely resemble the two main classes of utility-scale wind turbines on the market today: devices that run at very high speeds and which interface to the grid through full-converters, much like "type 4" wind turbines, and devices that run close to synchronous speed, but which have significant speed deviation from synchronous speed enabled by a double-fed topology that is much like "type 3" double-fed wind turbines.

Deployments of flywheels, as measured in both count of projects and total MW/MW-sec rating, are relatively small compared to other commercial storage technologies. Essentially all of today's commercially available flywheels are high-speed and use full converters. They are ideal for frequency regulation and other fast response services due to their fast response time and high power to energy ratio. High power flywheels can recharge in seconds[83]. They have very high efficiencies of over 90%[84]. Energy flywheels (as opposed to power flywheels) are being designed for longer durations with advanced technologies to reduce stand-by losses and increase round-trip efficiency. Cycle life is typically estimated at over 100,000 cycles[85], although newer configurations report 175,000 cycles[86]. There is virtually no wear-and-tear from cycling. The DOE Global Storage database lists 40 operational flywheel energy storage projects totaling 930 MW[87]. Twenty-one of these are used for frequency regulation applications. The following responses apply to these full-converter, high speed flywheels, unless otherwise noted.

How fast can it respond? In 2013, Sandia/EPRI/DOE reported for power flywheels: very fast response times ($T_{Activate}$+ $T_{ActivateFully}$) of 4 ms or less. Esave's flywheel solution claims a response time ($T_{Activate}$+ $T_{ActivateFully}$) in as little as 10 ms[88].

How much can it inject or absorb? 100% of capacity.

How long can it sustain this response? Typical FESS deployment has been designed to provide power for frequency regulation with sustain times on the order of 15 minutes. However, flywheels can also be designed to provide energy with longer durations, on the order

83 Ibid.

84 Ibid.

85 Ibid.

86 http://energystorage.org/energy-storage/technologies/flywheels

87 http://www.energystorageexchange.org/ accessed July 17, 2016.

88 http://www.esavecorp.com/?page_id=3873

of hours. An Amber Kinetics project in San Diego Gas & Electric is demonstrating a four hour flywheel[89].

Is there a particular ramp or shape to the response? No.

What determines each of these limits? Is there potential for it to be adapted? FESSs are incredibly fast. The technology response time, capacity that can be utilized, and sustain time are all suitable for AEMO needs. The constraining factor for FESS will be RoCoF detection and communication time.

Where, when and how have these capabilities been demonstrated?

In the US, Beacon Power installed a 20 MW (15 minute) flywheel in 2011 to provide frequency regulation to the NYISO market. Spindle Grid Regulation, who now owns Beacon Power, also installed a 20 MW (15 minute) flywheel in the PJM market to provide fast frequency regulation services. A 23 MW flywheel is in operation in Okinawa Power, providing frequency regulation[90].

A flywheel-battery hybrid with two 160kW flywheels and 240 kW (80kWh) of batteries was installed in Ireland by Schwungrad Energie to demonstrate various fast responses that could be provided in EirGrid's DS3 program. The power conversion system was designed for response times ($T_{Activate}$+ $T_{ActivateFully}$) of less than 20 ms[91] so that the system could provide a fast frequency response and limit RoCoF.

Industrial flywheels have long been used for specialized applications. For example, GE provided a 5MW, 40 ton flywheel on the dragline at the Usabelli Coal Mine in Alaska, installed in 1978[92]. It uses an early (silcomatic) power converter. It smoothes out the cyclic loading and regeneration of this mining load that had previously been highly disruptive to the rural grid to which it is attached.

The largest flywheels in the world are used in the high energy physics research communities to provide large amounts of power for very short periods (seconds) of time. For example, the Max-Planck Institute in Germany has been using a FESS since 1987 with three units that have capacities of 155 MW (9.7 second), 124 MW (4.4 second), and 108 MW (6.7 second)[93]. But these installations are not part of utility grid operation for discharge, as they are synchronous machines that depend on substantial deviations of terminal frequency to extract the energy from the rotor. (more below, in the next subsection).

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. There is no impact on

89 Akhil, A., Huff, G., Currier, A., Kaun, B., Rastler, D., Chen, S. Cotter, A., Bradshaw, D., Gauntlett, W., "DOE/EPRI 2013 Electricity Storage Handbook in Collaboration with NRECA", SAND2013-5131, July 2013.

90 Ibid.

91 http://www.energy-storage.news/news/flywheel-battery-hybrid-system-installed-in-ireland

92 http://energy-alaska.wikidot.com/usibelli-flywheel

93 http://www.energystorageexchange.org/projects/918 accessed July 17, 2016.

the flywheel from providing fast response or 100% response. Longer sustain may require design for energy, not power, or hybridization with a battery.

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? Flywheels have been around for a long time and are commercial, but they have not been widely deployed. The technology is moderately mature and there is potential for improvements especially as energy flywheels are further developed and increase the duration of the storage capacity.

How is the ability to deliver a fast frequency response service affected by nearby power system disturbances, such as transmission faults? Will the FFR service still be available in the period following a transmission fault? How could this be managed? This service should still be available following a transmission fault. The issue of DC bus stiffness discussed above applies to full inverter flywheels and double fed machines. The maintenance and recovery of DC bus voltage will depend on design, particularly control of the DC chopper for full converter flywheels. Good performance should be possible.

What is the potential for false or spurious triggering? How is this managed? As with the other inverter-based technologies in this section, this depends on the RoCoF detection scheme (see section 2.2.2.2).

Are there any other important considerations for the use of this technology to provide a fast frequency response service, to mitigate high RoCoF? No.

3.4.2 Flywheel Energy Storage Systems with Large Synchronous and Near-synchronous Machines

As noted above, there is a class of flywheels that use double-fed topology. These are discussed along with synchronous devices in this subsection. In general, this document is dedicated to discussion of technologies other than conventional synchronous generation and transmission alternatives to address the challenge of managing high RoCoF and severe frequency excursions. However, there is one emerging variation on traditional synchronous condensers that warrants discussion here: synchronous condensers with enhanced inertia. Most recently large (200MVA) class synchronous condensers with inertia constants of around 8 MW-sec/MVA have been built to support frequency, voltage and short circuit strength at the inverter end of a new HVDC project in Canada. This technology has several attributes that may be especially attractive for South Australia, and perhaps Tasmania as well. In addition to providing synchronous inertia, the devices will provide dynamic voltage support and short circuit strength. [Additional dynamic voltage support would have been welcome during the 28 September 2016 event.[18]] Capital costs for high inertia synchronous condensers are on the order of $25US/kW-sec.

There are also hybrid devices, that have been called "rotating stabilisers[94]", which operate as synchronous condensers down to the frequency nadir. At the frequency nadir, rather than

[94] "Rotating Stabiliser" has been trade-marked by GE.

reaccelerating, they switch to a full converter topology. This is accomplished by disconnecting the synchronous connection of the stator and feeding it through a full ac-dc-ac converter. In this second mode they share topology with the class of flywheel systems discussed above in Section 3.4.1. The rating of the converter is selected to match the desired power-time relationship needed for the application (more discussion in the next section). The inertia constant for these devices can be customized, and is reported to be available in range of 20-40 MW-sec/MW. This suggests that a single installation could maintain South Australia inertia at the low end present range of operation for conceivable future operation at zero synchronous generation. However, loss of the majority of the voltage support and short circuit contribution in this mode might be problematic.

While not exactly commercialized, there is potential to use the variable speed technology that enables variable-speed pumped storage hydro (VSPSH) as flywheel energy storage system. We have not discussed VSPSH in this section, with the understanding that longer term energy storage technologies are out of scope for this report. Further geographic topology for PSH is probably not good in South Australia, but the double fed technology for VSPSH is applicable to flywheels on a large scale. VSPSH is a small, but growing base of application. As far as we can determine, there are on the order of a dozen doubly-fed VSPSH projects built or under construction in the world. Ratings are mostly on the order of 100-200MW. The GE machines under construction are on the order of 200MVA machines.

3.4.2.1 Big Flywheel Discussion: Synchronous vs. Variable Speed

Some comparative discussion of large machines with high inertia is helpful.

For the sake of comparison, we introduce 200 MVA machine, with augmented inertia to an inertia constant H = 8 (at synchronous speed), which results in a total inertia of 1600MW-sec. This is a big enough inertia to have a noticeable impact on the RoCoF and frequency problems of concern in South Australia (and Tasmania).

The fundamental physics of a flywheel dictates that, to get power in or out, the speed needs to change. The energy stored is proportional to the square of the speed (duh!). A linear approximation is, therefore, that for each 1% change in speed the energy delivered (consumed) changes by 2%. A synchronous machine is, by definition, constrained to a speed proportional to the grid frequency. Therefore, the energy that can be extracted is constrained by the amplitude of the allowable frequency excursion. In laboratory settings, where synchronous flywheels are used to power plasma devices and particle accelerators, the machine is "spun up", and discharges its energy with a decaying AC frequency – desynchronized from the host grid. Interesting, but of no use to AEMO. For a synchronous machine in the NEM, the contribution of synchronous inertia (primarily) occurs within the frequency limits of the grid, i.e. before the UFLS triggers at 49Hz. Since 49Hz is a 2% deviation, 4% of the stored energy is useful: just like ALL the other synchronous generators that might still be running. The contribution is inherent, instantaneous and doesn't depend on fast sensing, communication or actuation. For our numerical example, the arresting energy delivered is:

MVA*H*((1)^2 – (1-dF)^2) = 200*8*((1.0^2)-(.98^2)) = 200*8*0.04 = 64 MW-sec.

With a double-fed arrangement, the machine speed can change over a much wider range. The enabling technology is the same as that of double—fed (type 3) wind turbine. The power-electronics provides excitation and enables slip. The converter rating is proportional to the slip speed. So, for our 200MVA, 1600MW-sec, and example 30MVA (nominal) ac-dc-ac inverter gives +/- 15% speed range, which means that the energy that can be actually used is 60% of the stored energy:

$MVA*H*((1+dSlip)^2 - (1-dSlip)^2) = 200*8*((1.15^2)-(.85^2)) = 200*8*0.6 = 960$ MW-sec.

Which equals the machine rated power delivered for 4.8 seconds: 15 times the energy of the synchronous machine. To a first approximation, the machine cost is the same (or somewhat higher), but the power electronics of rating about 15% of the machine, adds cost. With this arrangement, voltage control may be faster, and maintaining synchronism will be easier (this device will not be subject to traditional first swing stability constraints, as a synchronous device will necessarily be. But, aside from higher capital, with this arrangement there are other draw backs:

(a) the machine relies on control – so the 3 elements: measure, communicate, activate apply. This is *not* synchronous inertia, this is FFR.
(b) The machine does not contribute to short circuit strength, rather it exacerbates the low SCR (weak grid; high SNSP) problem.
(c) It is less widely established technology.

Nevertheless, the scale and agility of the device, coupled with large amount of arresting energy that could be delivered rapidly appear to make this technology attractive for this application.

The hybrid "rotating stabilizer" configuration is a significant improvement on some aspects of the trade-off between synchronous and variable speed. To continue the numerical example, while in synchronous mode, the device will deliver 4% of the stored inertia energy and all the other benefits of being a synchronous machine up to the time of the nadir. At the time of the nadir, the device switches to provide continued primary frequency response, by further decelerating the rotor. The converter allows for nearly all of the remaining rotational energy (1536 MW-sec) to be delivered to the system. The power rating of converter dictates the maximum power that can be delivered. So, for the numerical example, a 15% rated (i.e. 30MVA) converter, would allow 30MW to be delivered for about 50 seconds.

So, in comparison to the double-fed arrangement, which allows the full MW rating (of the machine) to be delivered, but which cannot extract all of the stored energy, this arrangement allows all the rotational energy to be delivered, but at a power level limited by the converter rating. Table 5 provides a concise comparison of the three topologies. The rotating stabilizer is normally synchronous, and switches to full-converter at the frequency nadir.

Table 5 Comparison of Available Energy for Flywheel Topologies

	Synchronous	Double-Fed	Full Converter
Usable Energy (% of ½ $I\omega_o^2$)	ΔF^2 (~4%)	$2\Delta\omega_{max}^2$: $\Delta s_{max} = Pe_{rated}$ (~60%)	Pe_{rated} = 100% (~100%)
SIR/SCR contribution	yes	no	no
FFR/PFR	No (not controlled)	Yes (Pmax = 100%)	Yes (Pmax = Pe_{rated})
Q	100% (machine Q rating)	100% (machine Q rating)	Pe_{rated} (converter Q rating)

Ultimately, it is unlikely that any of these devices will be justified solely on their inertia/FFR benefit, but rather they could be justified on all three attributes: inertia, dynamic voltage support and short-circuit contribution. As the discussion illuminates, there are a range of options available to the designer. These machines will *inevitably* be customized for the specific needs of the application, and trade-offs between performance in these 3 aspects, balanced with cost evaluation of no-load and full-load losses, and capital costs will have to be made based on detailed and specific requirements. These can only be given as the result of thorough systemic analysis. One design most emphatically does not "fit all" here.

3.5 Supercapacitor Energy Storage Systems

Supercapacitors, or ultra-capacitors, are electrochemical capacitors. They are similar to batteries in that ions migrate from one electrode to another through an electrolyte. Unlike batteries, there is no chemical reaction. Therefore wear-and-tear from cycling is not an issue and cycling does not impact lifetime. Electrodes are typically made of porous carbon which has a high surface area, which provides extremely high storage capacity. They have a much higher power density than other storage technologies discussed here. Because they have very low internal impedance, they can provide high currents. Key advantages are high cycle life (100,000 cycles), high power, high efficiency, and fast response times. Supercapacitors can be charged and discharged rapidly, in 1-10 seconds.

Disadvantages include low energy, cost, and need for power conditioning. Power conditioning is necessary because capacitor voltage is proportional to the stored charge in the system. Supercapacitors are used in the transportation industry for regenerative braking, voltage drop compensation, or for short bursts of power, and in industrial applications as uninterruptible power supply and also in consumer electronics. Their use in the utility sector is relatively new, and they tend to be used in hybrid applications with batteries or as part of advanced lead-acid batteries.

How fast can it respond? SBC quotes response times ($T_{Activate}$+ $T_{ActivateFully}$) of 10-20ms[95].

How much can it inject or absorb? This would depend on the power conditioning. Nearly 100% of capacity can be used if power conditioning is designed for this.

How long can it sustain this response? A disadvantage of supercapacitors is the short sustain time, on the order of seconds to minutes.

Is there a particular ramp or shape to the response? There is a shape to capacitor response so power conditioning is needed to form the response.

What determines each of these limits? Is there potential for it to be adapted? Hybrid technologies, such as advanced batteries that combine conventional batteries with supercapacitors, can ensure a very fast response that can be sustained for minutes to hours.

Where, when and how have these capabilities been demonstrated? The DOE Energy Storage Database lists 28 projects with a total of 76 MW capacity that have been deployed, mainly to provide transportation services and voltage support. They have also been deployed in their use in advanced batteries.

In 2014, ENDESA in Spain commissioned a 4 MW (6 second) supercapacitor system to stabilize voltage in La Palma in the Canary Islands. The system helps to avoid underfrequency load shedding in a contingency event.

In Palmdale, CA, a 450kW supercapacitor from Maxwell Technologies provides 30 seconds of storage as an uninterruptible power supply for a water treatment facility[96].

At University of California in San Diego, a 28 kW supercapacitor system provides smoothing and firming for a concentrated PV system[97].

In early 2016, Duke Energy installed a hybrid energy storage system with a 277kW (8kWh for a duration of nearly 2 minutes) supercapacitor[98,99] and a 50kW (200kWh) battery that provides fast solar smoothing and longer-term load shifting. Fast, high power responses are provided by the supercapacitor to reduce thermal stress of the battery and not degrade battery lifetime.

---

95 SBC Energy Institute, “Electricity Storage”, September 2013.

96 Ibid.

97 Torre, W., Borneo, D., and Washom, B. “Energy Storage for Integration of Renewable Generation at the University of California – San Diego”, Electrical Energy Storage Applications and Technologies, San Diego, CA Oct 20-23, 2013.

98 https://news.duke-energy.com/releases/duke-energy-to-put-new-battery-and-supercapacitor-system-to-the-test-in-n-c

99 http://wininertia.es/projects/rankin/

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. There are virtually no wear-and-tear costs for supercapacitors, so operating in this manner is not a problem. O&M is very low. Cycle life is typically estimated at 100,000 cycles. One manufacturer estimates 1 million cycles with 20% reduction in rated capacitance[100].

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? This technology is commercially available and is based on the mature battery and capacitor industries. However, supercapacitors have not been deployed much for utility-sector applications and therefore, there is likely to be some potential for reduced costs and improvements in power conditioning with widescale deployment.

What is the potential for false or spurious triggering? How is this managed? As with the other inverter-based technologies in this section, this depends on the RoCoF detection scheme (see section 2.2.2.2).

Are there any other important considerations for the use of this technology to provide a fast frequency response service, to mitigate high RoCoF? For utility sector applications, this should not be an issue but self-discharge rates are reported to be ~5-40%[101]. Supercapacitors are adversely affected by high temperatures but this should be manageable in system design. This does mean that overheating can result if an supercapacitor were exposed to a ripple current, but RoCoF applications should not cause this issue.

3.6 Solar PV

Solar Photovoltaic (PV) generation is enjoying a worldwide explosion in growth. The two most basic components of PV for provision of AC power are common to all applications, in power ratings ranging from a fraction of a watt to utility scale projects approaching (or even exceeding) a Gigawatt. The first component is a collection of photosensitive semiconductor cells that when subjected to photons of appropriate wavelength, produce energetic electrons that will flow, when allowed to do so: that is, they produce direct voltage and current. The product of this DC voltage and current is power, which must be converted to 50Hz AC power by inverters. PV installations consist of one or more of these modules, normally connected in parallel. For projects of utility scale, AC collector systems (sometimes 'reticulation' in Australia) resemble those of wind plant, with a supervisory control providing intelligent interface between the grid and the farm. The most basic concept is shown in Figure 29.

100 http://www.maxwell.com/

101 Luo, X., Wang, J., Dooner, M., and Clarke, J. "Overview of current development in electrical energy storage technologies and the application potential in power system operation," *Applied Energy* 137 (2015) 511-536.

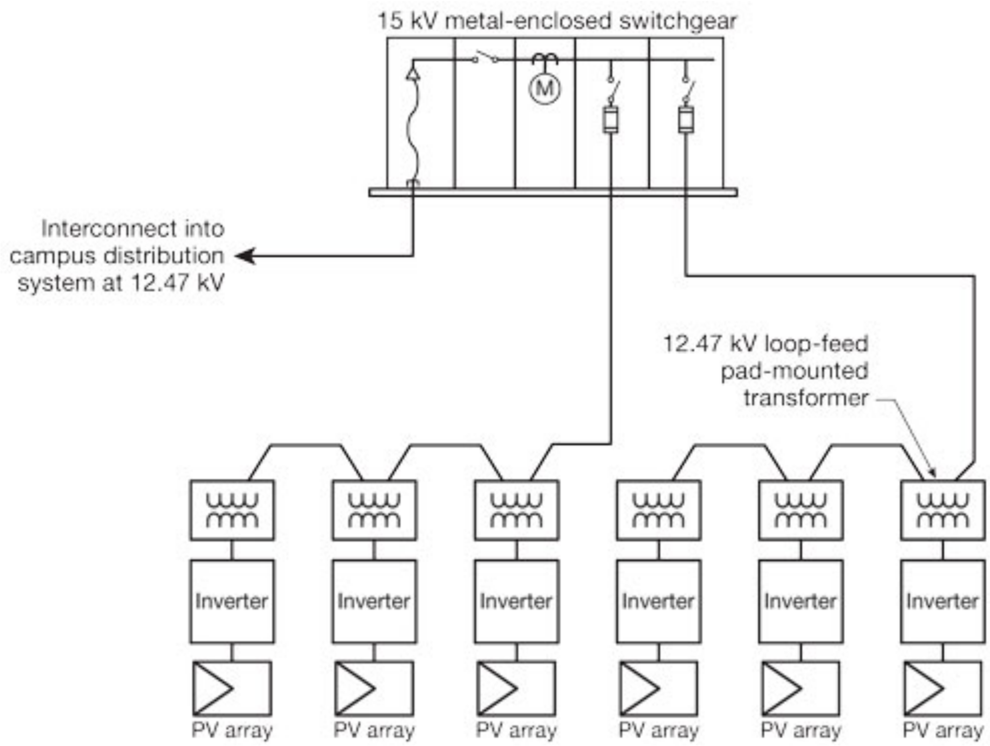

Figure 29 Basic PV module[102]

3.6.1 Solar PV Components and FFR

For this discussion of possible fast frequency response, a few physical elements of PV are critical:

1. DC Power Rating. Individual photocells typically have voltage rating on the order of 2V, and current rating very roughly proportional to the area of the cell. The details of voltage, rating, sensitive wavelength, temperature sensitivity, etc. vary with different cell designs and materials, and are not particularly important here. The DC rating of the collection of cells (i.e. a "panel") dedicated to a specific inverter is important. The DC rating (for this discussion) is the maximum DC power that can be produced under conditions of maximum insolation (sunlight energy intensity at the panel). DC rating is independent of the inverter.
2. AC Power Rating. The inverter serves the function of converting the DC power to AC. The cost of the inverter is dominated by the AC current rating, although the rating is typically given in kVA at nominal AC voltage. The key point is that there is no

[102] www.solarprofessional.com

fundamental requirement that rating of the inverter "match" the DC rating of the panel. In practice, this has significant implications for FFR, as will be discussed below.

3. Tracking. The orientation of the panel relative to the sun dictates how much of the available insolation energy is converted. In general, PV installations can be fixed, single axis trackers, or dual axis trackers. Tracking installations physically orient the panels so that they capture more energy over the course of a daily solar transit. For a given day and DC rating, the DC power available from the 3 types of tracking (for similar DC rating) are conceptually shown in Figure 30.

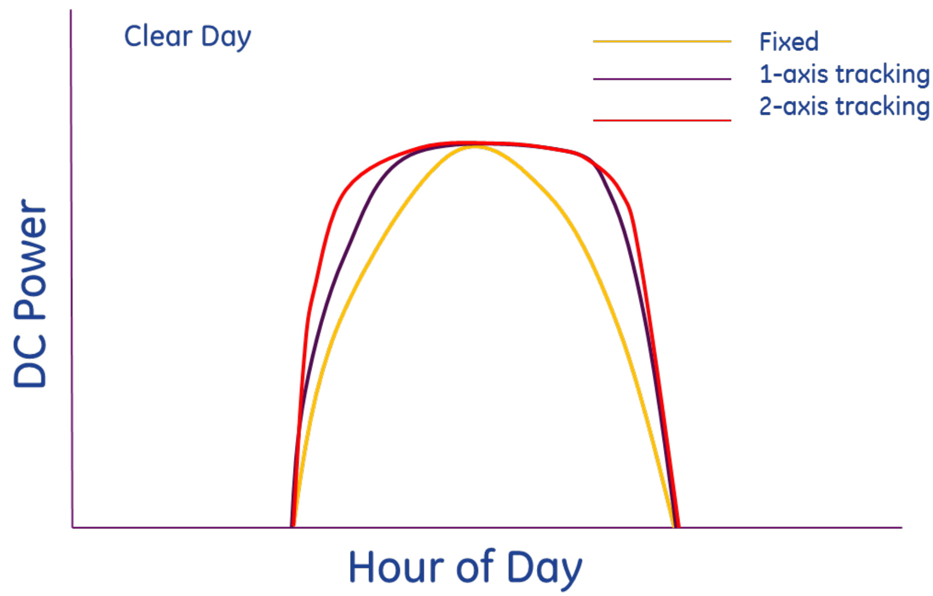

Figure 30 Illustration of DC Power Impact of Tracking

4. Time of Year. The impact of time of year is illustrated in Figure 31. The shorter days and lower incidence angle affect the energy production. This seasonal variation has some importance, relative to the inverter rating, as will be discussed below.

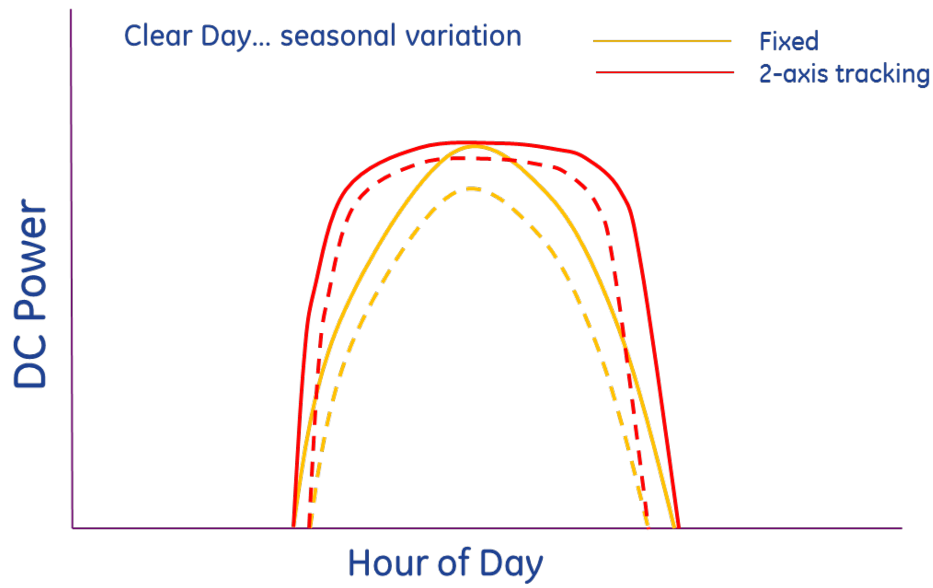

Figure 31 Illustration of Time of Year Impact

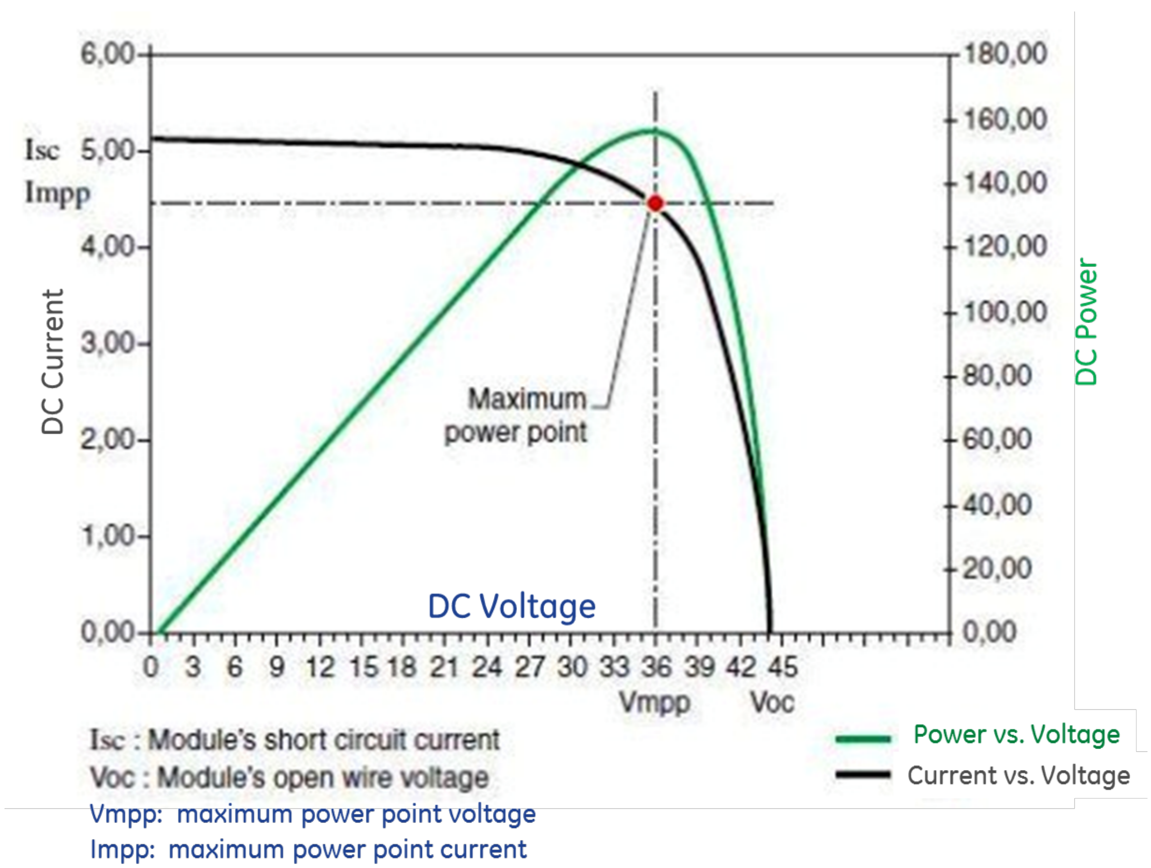

Figure 32 Maximum Power Point Tracking Illustration

5. Maximum power point tracking (MPPT) control. The amount energetic electrons liberated from the photovoltaic cell is sensitive to the voltage across the cell, the amount of insolation and the temperature of the cell. In **Figure 32** the relationship between DC voltage, current and power is illustrated for one level of insolation and cell temperature. The insolation and temperature of the cell vary continuously, with insolation moving quite rapidly under some cloud conditions. One of the functions of the inverter is to maintain the best DC voltage for the present condition, by means of a maximum power point tracking (MPPT) control. The MPPT is the means by which the solar PV maintains the minimum of: maximum available DC power, maximum AC power, i.e. the inverter power limit; power setpoint, i.e. a curtailment. The MPPT is normally a slower closed loop control than the inverter current control loop and the supervisory voltage control loop. In order to maintain control stability between these controls, their speed of response is typically separated by an order of magnitude. Consequently, the MPPT control has a time constant on the order of a second.
5. Supervisory control. For utility scale PV installations (i.e. those which are most likely to be able to provide FFR), there is typically supervisory control that has many of the same functions as a wind plant supervisory control. The supervisory control will accept curtailment (and other) instructions from the host utility, and communicate them to the individual inverters for implementation. As in wind plants, there are latencies in control and communication paths that limit how fast the supervisory control can transfer commands to the individual inverters. And again, like with the wind plants, the industry has not yet been required to drive these latencies to very low levels.

3.6.2 Curtailment, Overload and Provision of FFR

Unlike wind, there is no inherent energy storage available via kinetics (i.e. the rotating mass). Therefore, if a PV resource is to deliver any fast boost in power output, the inverter must be operating at an AC power transfer level that is LOWER than the available DC power from the PV panel at that point in time. Under conditions when the AC inverter (and balance of plant) is *not* the limiting condition, this requires that the power production be curtailed to a level below the available power. This is a nearly perfect analog to the pre-curtailment of the wind power necessary to provide sustained FFR or PFR, as described above in the wind section. Note that it is possible to have (battery or supercapacitor) energy storage integrated with the solar PV, but aside from possible savings associated with shared infrastructure (inverters, transformers, balance of plant), there is no intrinsic *performance* synergy associated with co-location (or hybridization) of the energy storage with PV. Otherwise, the energy storage discussion in the previous sections applies.

The steady-state and short term overload rating of the PV inverter can play an important role here. While it is natural to think of the rating of the inverter in terms of kW, or kVA, the physical reality is that the rating is dominated by the *current* carrying capability of the semi-conductor valves. A simple illustration of the concept is shown in Figure 33. This figure is greatly simplified, but captures the predominant character of the inverters not only for PV, but also for type 4 wind turbines, and for essentially all inverter-based energy storage devices. The figure is for a four-quadrant device (i.e. power can come in from the grid on the left – obviously only meaningful for energy storage and loads), and provide reactive power to the grid (northern hemisphere) or consume it (southern hemisphere). A key point is that P (active power) and Q (reactive power) are controlled independently *as long as the device is not at its limit*. But when the combination of P and Q demand drives the operating point against the circle, the control must decide which takes priority. The issue is further complicated by the reality that these limits are really *current* limits: the figure as labeled is meaningful for nominal terminal voltage. But when the voltage rises, the circles get bigger, and when it sags the circles shrink. (There are several additional nuances here, but they are not very important for this discussion). Finally, the outer circle, which represents overload capability, is complex. The degree of overload that can be tolerated is a function of design and time. Depending on the inverter design, this overload capability may be very small and very short…i.e. a few percent overload for fractions of seconds. But other designs may have more capability, either as a result of more robust design for other considerations, or as a deliberate incremental capability for short-term overload.

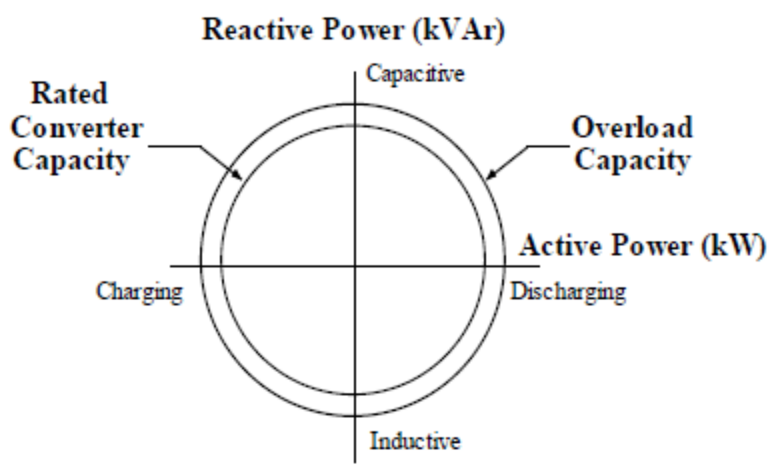

Figure 33 Inverter Rating Concept

Overall, this trade-off between active and reactive current capability, and system voltage, can be important for FFR. Two, *not* mutually exclusive, options exist:

1. **Give active power priority over reactive power**. This especially makes sense if the inverter is operating *under-excited* (i.e. consuming reactive power). It will raise the voltage, allowing more power to be delivered for a given current. Conversely, when the inverter is supplying reactive power, reducing the VARs to make room for active power will decrease the voltage, increasing the current necessary to deliver the same active power. It will be systemically dependent, and it is possible that the voltage reduction will negate the benefit or create other systemic problems. These considerations are situational, and will tend to be greater in weak grid situations.

2. **Drive the inverter into the overload range.** This must be done carefully, and with knowledge of the design capability of the inverter. Considerations such as ambient temperature and pre-disturbance valve junction temperature might be important, if squeezing every last bit of power out is important.

These points *only make sense* if it is the inverter that is limiting the delivery of active power. The inverter cannot cause the DC panel to make more power than the maximum allowed by the instantaneous insolation.

The interplay between these two options, the system requirements and the available DC power can be rather confusing. Phasor diagrams help illuminate the relationships. **Figure 34** shows the current for a normal operating condition (blue dot) in which the current limits of the inverter are not active. The blue phasor, shows active current (I_p) being injected to the grid at a particular level of DC power (insolation), given by the vertical dotted blue line. The reactive current, I_q, in this quadrant is being injected to the grid to support voltage (presumably in response to required closed-loop voltage regulation). As long as the total current, I_{total}, is within the rating circle, the active and reactive injections are independently controllable.

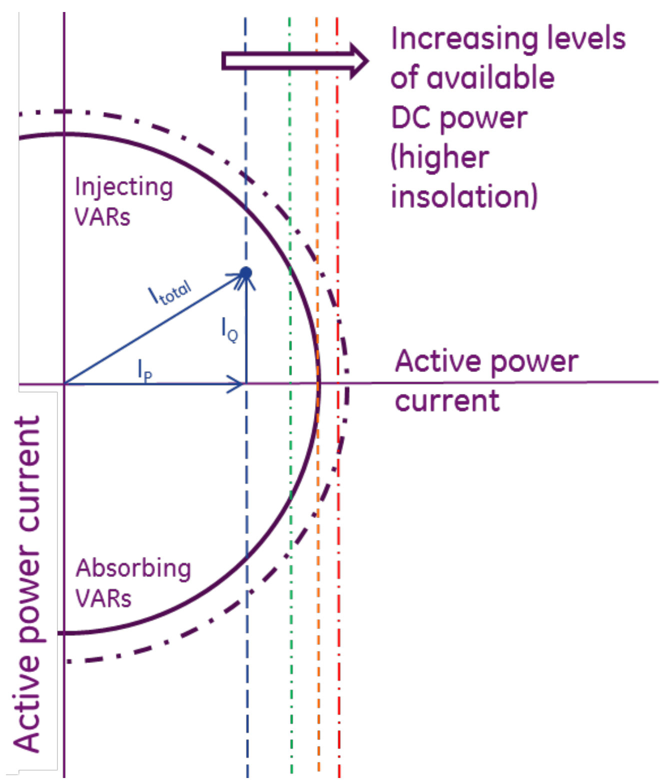

Figure 34 PV Current Phasor for Normal Operation

As available DC power increases, the operation moves towards the steady-state current limit. In **Figure 35**, the DC power is assumed to have increased to the green vertical dashed line, while the reactive current requirement remains the same. This new operating point, the green dot, represents operation at the steady-state current limits of the inverter. In the event that the system reactive power demands (for I_q) *and* the available DC power (I_p at this voltage) together *exceed* the current rating, then the inverter must establish which function takes priority. If reactive support takes priority, which is typically the case *up to the power factor required by interconnection requirements,* then active power will be limited. Alternatively, reactive current can be reduced to *make room* for active current injection.

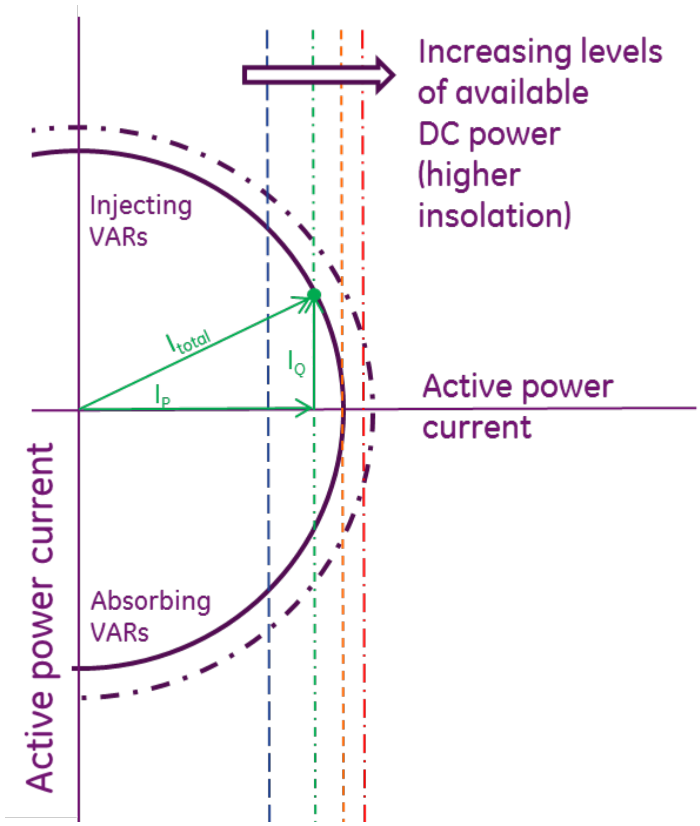

Figure 35 Normal Operation at Steady-state Current Limits

The concept of giving priority to active power for high levels of insolation is shown in Figure 36. Here, an even higher level of insolation, given by the orange vertical dashed line, results in a reduction of reactive current injection. In the limit, the inverter could go to unity power factor, without exceeding the steady-state current rating of the inverter.

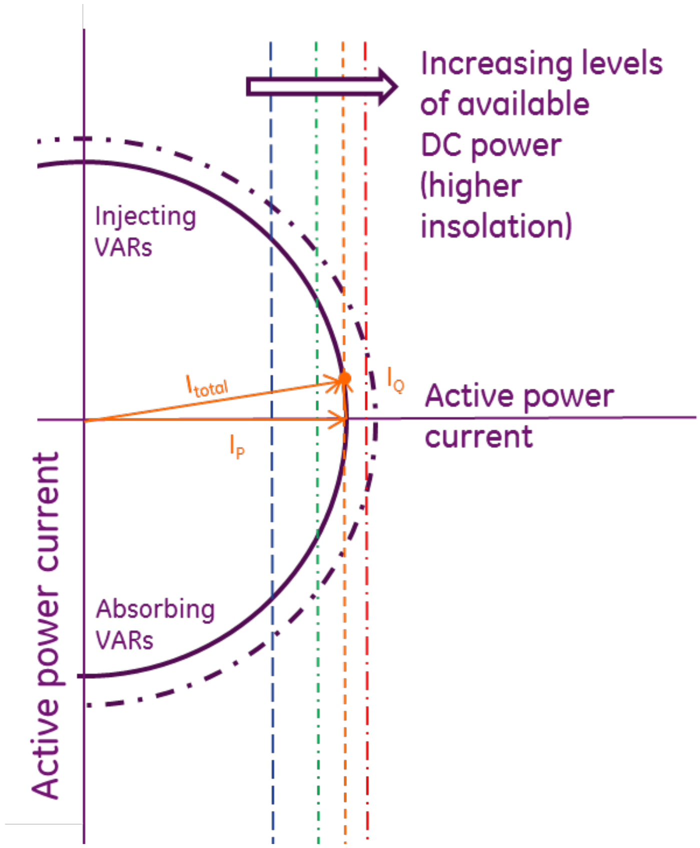

Figure 36 Active Power given priority within Steady-state Rating

The concept of taking advantage of possible short-term overload rating of the inverter, as illustrated by the dashed purple circle, is shown in Figure 37. The short-time overload rating circle can be time, temperature and operating history dependent. The two concepts, giving P priority and using short-term overload capability, can work together to provide control options for achieving FFR.

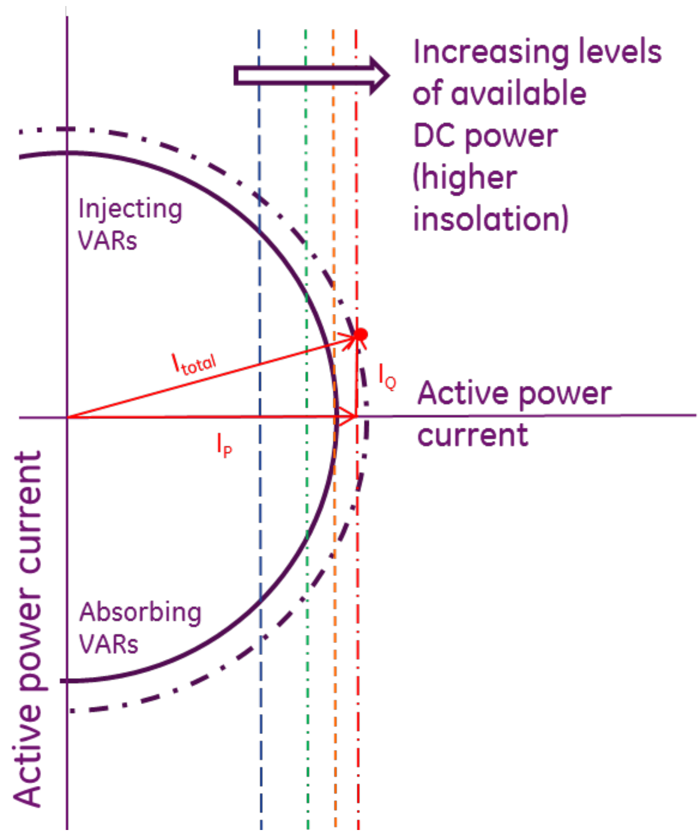

Figure 37 Operation with Short-time Overload of Inverter

These current-rating and DC power relationships raise an interesting new option for FFR, as discussed in the next section.

3.6.3 Inverter and Panel Rating

Historically, the biggest cost component in photovoltaic systems was the semi-conductor panel. Further, until recently, the overall cost of PV energy was considerably higher than from competing generation resources. Consequently, the economics of PV design tended to strike a balance between using the absolute least expensive components for the rest of the system (i.e. the balance-of-plant), including the inverter, and making sure that every kW that the panels might produce would be delivered to the user (meter). Therefore, inverters were commonly rated to meet the maximum DC power of the panel. Further, the inverters were assumed to operate at unity power-factor, thereby eliminating any extra current rating (and cost) associated with delivery of reactive power.

Recently two important changes have developed: the cost of panels has dropped precipitously, making the relative cost of inverters in the overall system much larger, and grid authorities are demanding (through grid codes, etc.) that PV provide voltage support and have reactive power capability.

These factors drive different design decisions that are important for this discussion of FFR capability. Specifically, it is becoming common for the steady-state active power rating of the inverters to be *less* than the maximum DC power of the panels to which they are matched. The concept is illustrated in Figure 38.

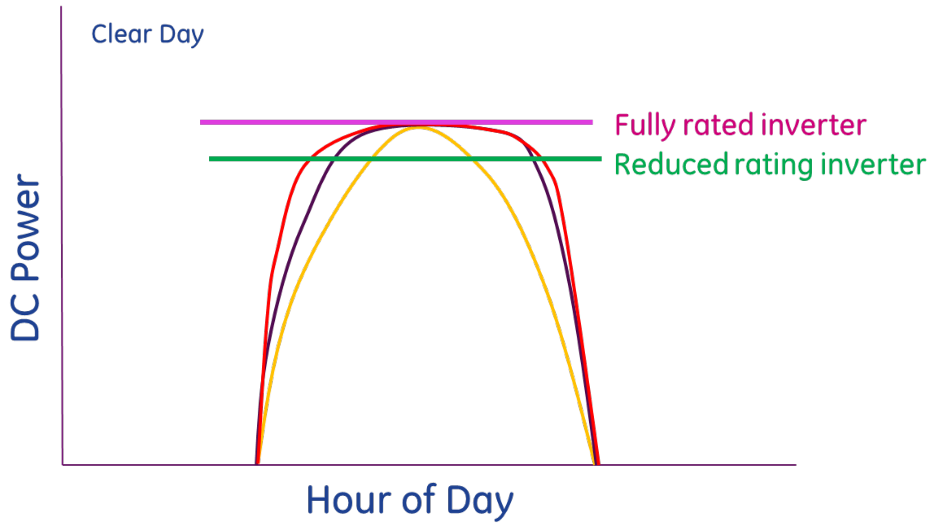

Figure 38 DC power (different trackers) vs. AC Inverter Rating

The implication of this rating difference for provision of FFR is potentially profound. If the under-rated inverter is designed to have a degree of short-term overload *and/or* the inverter is given active power priority over reactive power delivery, then there is extra active power available for delivery to the grid. The concept is illustrated in Figure 39.

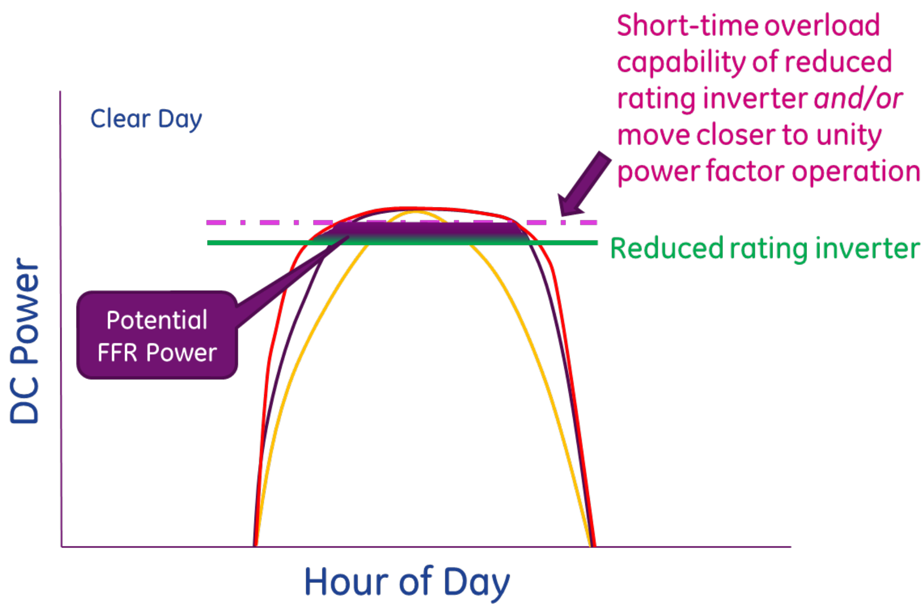

Figure 39 FFR capability for reduced rating PV inverters

This means that during any period in which the insolation exceeds the steady-state capability of the inverter, there will be essentially *zero* opportunity cost to the PV plant to provide FFR up to the short-time limit of the inverter and the instantaneous insolation. The owner of the PV *will incur* capital costs to have this capability. Further, this is new ground. There is not, at this writing, industry precedence for this approach.

3.6.4 PV Experience with Fast Control Response

Recently NREL conducted extensive testing on a 20MW utility-scale PV plant in Puerto Rico[103]. The plant, which was developed by FirstSolar and is owned by AES, was subjected to a number of tests (as quoted directly from the NREL report):

> "This demonstration showed how active power controls can leverage PV's value from being simply a variable energy resource to providing additional ancillary services that range from spinning reserves, load following, ramping, frequency response, variability smoothing and frequency regulation, to power quality. Specifically, the tests conducted included variability smoothing through automatic generation control, frequency regulation for fast response and droop response, and power quality."

Results for the fast frequency response tests are shown in Figure 40. Again, to quote the NREL report:

[103] www.nrel.gov/pv/news/2016/21630.html

”There was a delay of approximately 50-100 milliseconds between the frequency change and the beginning of the FFR by the plant. Implementing sophisticated control theory and fine tuning the power and current gains could help with minimizing the oscillations shown in the traces below. In general, more care is needed to fine tune and set controls to achieve better long-term performance. But with respect to the speed and power increase of the frequency response, the tests were successful.“

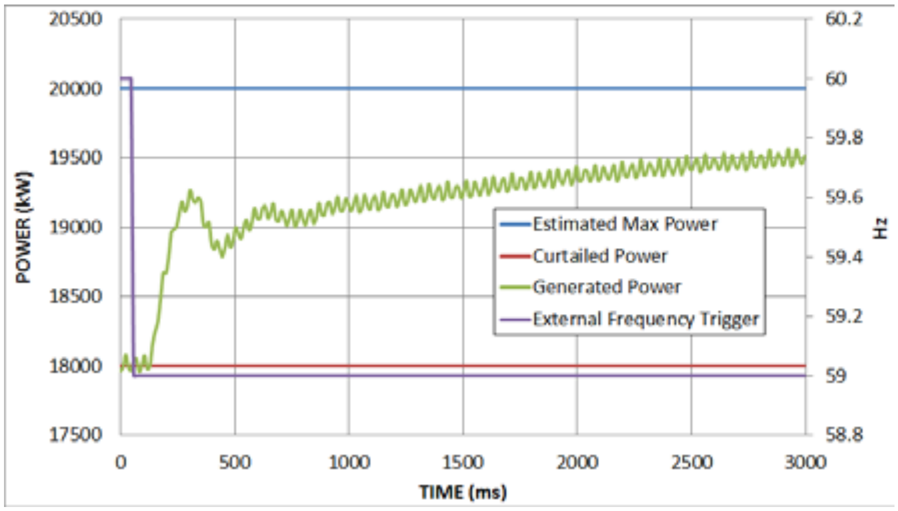

Figure 40 FFR Tests for 20MW PV Plant[104]

Limits on speed of response. The response here (Figure 40) appears is on a similar scale of speed to the BESS results presented above. The physics of the converter are nearly identical between the two, but the character of the DC link is rather different and can have an impact on the speed *and amplitude* of the response. Specifically, the DC voltage on the terminals of a battery string is relatively “stiff” with a significant internal resistance. The resistance creates a drop in the DC voltage when the discharge current jumps up (and conversely a rise, during a jump up in charging current). Consideration of this droop is included in the current control, which adjusts pulse width modulation (PWM) firing accordingly. The DC voltage for a PV panel is less stiff, and itself relies on the response of the photosensitive semi-conductors. The net result is that manipulation of the active current drawn from the panel has a substantive impact on the DC voltage, which must be managed by the converter. This has the potential to impact the speed of response. To-date, PV has not been required in commercial operations to adjust active power in response to e*xternal* instructions at extremely high speeds.

The fast response using pre-curtailment of the PV system overrides the MPPT control and operates *off* of the PV module’s maximum power point. The inverters operate on the over-voltage side of the MPPT curve shown in Figure 32, and when FFR action is needed, the MPPT nearly instantaneously shifts the voltage to the maximum power point, providing an increase

104 Advanced Grid-Friendly Controls Demonstration Project for Utility-Scale PV Power Plants; V. Gevorgian, B O’Neill, NREL/TP-5D00-65368, January 2016.

in power output. The DC voltage dips along with the response, and implications on system stability need to be considered. The response time of this approach depends on the equivalent input capacitance of the PV, which is small. However, recent tests on present full sized (300MW) state-of-the-art utility scale PV projects suggest that response times of 100ms should be doable. Further, PV inverter experts interviewed say it should be possible to do this within 50ms of receiving an instruction.

Control Philosophy. The relative benefits and disadvantages of open-loop vs. closed-loop FFR controls that were described in detail in the wind section apply to solar PV. However, this discussion and debate has, for practical purposes, not begun in the industry.

3.6.5 Costs

There is not yet sufficient commercial experience to determine capital costs for FFR (or PFR) in PV plants. Because FFR affects inverter software controls to a greater extent than physical hardware (or inverter rating, which would impact copper and silicon content), the incremental cost is thought to be modest. Historically, inverter costs have been a small fraction of the overall PV plant cost. But, per the discussion in Section 3.6.3, this relationship in component costs is changing. NREL's Q1 2015 benchmarks for breakdowns of PV plant costs find that inverters cost $0.29/W out of a total installed cost of $3.09/W for residential, rooftop PV, or about 9% of the total cost[105]. If, hypothetically, PFR were to increase inverter costs by 25%, the total rooftop PV system cost would increase by 2%. For utility-scale PV systems, inverters cost $0.11/W out of a total installed cost of $1.77/W, or about 6% of the total cost[106]. A hypothetical 25% increase in inverter costs would increase the total utility-scale PV system cost by less than 2%. Again, as was noted with the wind discussion, a 1% to 2% increase in capital cost can represent a substantial penalty against the margins or return on investment for a project. This is especially true if the increased capital costs are not factored in while the project is being designed and financed. Retroactive requirements are likely to be much more expensive.

3.6.6 Synopsis of Solar PV FFR

How fast can it respond? Response times on the order of 100-200ms have been demonstrated.

How much can it inject or absorb?[107]. Up to the short-time overload limit or maximum insolation.

105 Chung, D., Davidson, C., Fu, R., Ardani, K., and Margolis, R. "U.S. Photovoltaic Prices and Cost Breakdowns: Q1 2015 Benchmarks for Residential, Commercial, and Utility-Scale Systems," NREL/TP-6A20-64746, Sept. 2015. http://www.nrel.gov/docs/fy15osti/64746.pdf

106 Ibid.

107 HDR study for PacifiCorp, July 2014

How long can it sustain this response?. Up to the short-time overload time limit, if converter limited, or sustained longer at the maximum insolation if the inverter is within its continuous rating.

Is there a particular ramp or shape to the response? Trapezoidal behavior should be possible. Closed-loop control response should be possible.

What determines each of these limits? Is there potential for it to be adapted?. There are no torsional or mechanical limits; speed is primarily a function of control limitations, including coordination with MPPT.

Where, when and how have these capabilities been demonstrated? Very limited. The NREL tests quoted here were performed on a commercial installation. Several large PV plants in the US are equipped with a full spectrum of voltage/reactive power controls, and accept curtailment orders. We are unaware of any commercial installations that presently offer frequency control.

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. For overload, as described above, there should be no marginal cost. For curtailment, the opportunity cost is that of the solar energy (and lost tax benefits/subsidies, if any).

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? Not mature now. Per detailed discussion above, potential for development seems highly promising.

How is the ability to deliver a fast frequency response service affected by nearby power system disturbances, such as transmission faults? Will the FFR service still be available in the period following a transmission fault? How could this be managed? PV tolerance to faults and high RoCoF can be good, but practice in some places has caused distributed PV to have poor grid fault and frequency tolerance. The consideration of DC bus stiffness (as discussed above) can be a contributing factor in the recovery following a fault. Maybe PV systems completely block (no current) during voltage depressions and faults.

3.7 Load Based Resources

Load resources can provide high quality FFR. For an under-frequency, negative RoCoF event, instead of asking a generator to rapidly increase power, we may elicit a faster response by asking load to drop. For example, in ERCOT, loads provide half of the interconnected power system's Frequency Responsive Reserves that support the system in a sudden loss of generation event. ERCOT's FFR triggers at 59.7 Hz or 59.8 Hz and provides full response in 500 ms. Loads that provide fast response have under-frequency relays that drop the load when a preset threshold is exceeded. Because loads, storage and switchable devices (e.g. dynamic brakes) can provide FFR, it's important to ensure that all resources are allowed to participate in potential markets for these services. These resources may prove to be more effective and economic than PFR provided by generation resources in maintaining system security.

For AEMO, participation by loads in an SPS that detects separation and other conditions which warrant controlled load shedding (e.g. import overload that risks separation) makes a great deal of sense. This option is integral to the extended analytical illustration provided below in Section 6.

3.7.1 Inertia Considerations

One of the considerations in using load as an FFR resource is that load may also be providing inertia, which is helping to mitigate RoCoF. Therefore, it may be necessary to assess the characteristics of potential load-based resources and accept those that can provide the greatest net benefit to the system. The power industry is struggling with development of good data and understanding of the inertial contribution of loads for large grid events. As has been noted before, the extreme nature of the problem facing South Australia makes this question nearly uncharted waters.

3.7.2 Tripping vs. Blocking vs. Continuous Control

Modern commercial and industrial loads often have a substantial component of electronically controlled motors and processes. This gives rise to a class of load participation in providing frequency response that is more sophisticated and more valuable than traditional breaker switched load shedding. As with the converters that control most of the other technologies discussed here, the front-end converters for many loads can be temporarily blocked or reduced in power levels in response to an appropriate signal. That signal could be locally measured frequency or a trigger (as from an SPS). The economics of using load this way varies a great deal. Considerations include the cost to the provider of interrupted or reduced throughput, capital costs, restart costs, to name a few. Economies of scale very much apply here. In the US, one of the more active participants in providing this service are aluminum smelting "pot lines", which can interrupt power consumption very rapidly and with little economic consequence if the duration is kept within bounds. Other loads, such as heating and cooling, tend to be good candidates.

3.7.3 Highly Distributed Responsive Loads

There is a practical limit to how much industrial load may be available to participate in fast voluntary load shedding. Therefore, it may be advantageous to consider means to get FFR from highly distributed loads. Rapid communication to a larger number of load nodes appears to be gaining practicality, as the Arizona example cited in Section 2.2.2.1 describes.

Another option maybe autonomous load controls. For example, in the UK, the National Grid Company (NGC) has introduced a special tariff which rewards large-scale consumers who agree to provide a limited form of frequency response. Called "Frequency Response by Demand Management" (FCDM), the scheme involves the placing of certain large loads behind frequency-sensitive relays which isolate the load when the frequency falls below a pre-set level, often 49.7 Hz[108]. Pacific Northwest National Laboratory in the US is promoting "Grid Friendly Appliances" which are similar to dynamic demand appliances being promoted in the UK. The use of highly distributed loads to provide FFR is likely to be subject to the constraints and uncertainties of rapidly measuring RoCoF and frequency as enumerated above. Waiting until frequency drops by (say) 300mHz in South Australia uses up a significant amount of the available time for FFR. Information on how quickly small domestic devices (e.g. water heaters, pool pumps) might respond to stop/block signals is limited, but the underlying physics suggests that rapid current interruption should be possible. Hydro Quebec has a pilot project for frequency sensitive control of hot water heaters. They aim to trigger within 250ms. For very fast actuation, again the pressing issue is likely to be mis-triggering associated with having frequency or RoCoF triggers too tight. Many classes of loads will tolerate occasional interruption, but will be unwilling to tolerate (for example) multiple interruptions per day.

To get higher fidelity trip signals, variations on transfer trip are required. The types of communication techniques discussed in Section 2.2.2.1 (and installed at Salt River Project, Arizona in the US) are suitable for rapid communication of trip signals to a large population of activating nodes. The challenge is likely to be identification of enough activating loads, and justification of the well-developed fiber-optic communication necessary to get the speed of response required in the South Australia situation.

One of the significant challenges associated with tripping highly distributed load is to know how much load will be disconnected at any given time. The activation of load blocking is, by its nature, discrete. The load is either blocked or not. Thus, droop or proportional response requires sophisticated distribution and time staggering across many binary triggering signals. If this *embedded* UFLS response represents a modest portion of the total response, tuning or differentiation is less important. But when this approach represents the majority of the response, then direct measurements, state-estimators, or other schemes to determine the amount of load that will be tripped *at any present* instant are needed. The technology exists, but the data gathering and computational burden for this approach are not trivial.

3.7.4 Parallels to Industrial Facilities

The situation faced by South Australia has strong parallels to problems occasionally faced by large industrial facilities that generate some of their own electricity, and purchase some from the host utility grid. GE has experience in design and evaluation of schemes for systems that allow "trip to island". As with South Australia, in situations where opening the connection from the facility to the grid results in excessive power-load unbalance, tripping of some loads is required in order to maintain stability. It is common for industrial facilities (e.g. chemical, refineries, manufacturing,..) to have a range of sensitivity for their loads, from highly sensitive (i.e. must not trip due to safety or high costs) to low sensitivity (i.e. easily tolerates occasional interruption). Differentiation between these loads makes it difficult to send "simple" breaker open signals, since that can result in critical loads being interrupted. Under such

108 http://www.dynamicdemand.co.uk/current_work.htm

circumstances, it is possible to use (a) signals passed through motor-control centers, that can differentiate the loads, and shut-down or back down those that can be tolerated, or (b) to place RoCoF or UFLS relays directly on the interruptible loads. Again, the primary question is one of speed for either option, and a further question for the local relay approach is the consequence of false triggers.

3.7.5 Synopsis of Load Based FFR

How fast can it respond? Communication to substations (T_{Signal}) is possible within 2-3 cycles. Blocking ($T_{Activate}$+ $T_{ActivateFully}$) depends on the load to be interrupted.

How much can it inject or absorb? This will be dependent on the sophistication of the state estimator (I.e., the ability to continuously and accurately track the amount of load that can be tripped by the FFR scheme at any given time).

How long can it sustain this response? Depends on the blocking mechanism, but assuming that relief is sustained until released by the grid operator is a good default.

Is there a particular ramp or shape to the response? Response is step-wise.

What determines each of these limits? Is there potential for it to be adapted? Speed of detection; speed of blocking.

Where, when and how have these capabilities been demonstrated?. Examples of fast load tripping have been implemented at Salt River Project, Arizona and Southern California Edison in the USA. See Section 2.2.2.1.

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. Customers will estimate opportunity and lost-revenue costs, based on specifics of load interrupted.

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe?. This is well established art, in slower time frames. The primary issue and potential for advancement (and application to South Australia) is whether time delays and latencies can be significantly reduced, without creating unacceptable risk/frequency of mis-operation. With the right economic signal, this technology could advance substantially.

How is the ability to deliver a fast frequency response service affected by nearby power system disturbances, such as transmission faults? Will the FFR service still be available in the period following a transmission fault? How could this be managed? This is primarily an issue of mis-operation, which depends on how and where the trip command is generated. If the trigger is directly initiated by the opening of the AC intertie, detection is not sensitive to post-fault distortions in voltage/current waveforms. And provided the trigger signal can be quickly and securely transmitted to the loads being tripped, FFR action should be very dependable. However, if the trigger signal is determined locally (near the load) by a frequency or RoCoF transducer, then there is some risk that distorted waveforms or local dynamic effects during and immediately following a fault could cause mis-operations (i.e., falsely triggering when not needed or failure to trigger when needed).

3.8 HVDC Transmission

This discussion focuses on the ability of HVDC to contribute to FFR during the period before the frequency nadir (i.e., the frequency arresting capability of HVDC).

3.8.1 Physics

HVDC power transfer is regulated by power electronics and programmable controls which are able to respond in a small fraction of a second. This capability can be used in a few different ways to respond to a power-load imbalance (either underfrequency or overfrequency) in an islanded system.

Fast ramp down in HVDC power (runback). This is a commonly applied control feature where power is rapidly reduced in response to an overfrequency event. For Voltage Source Converter (VSC) systems, AC bus voltage is regulated by the HVDC control system. For conventional line-commutated converter (LCC) systems, power reductions can produce AC overvoltages if there is an excess of reactive compensation, so some reactive power banks may need to be tripped. This is mature technology.

Fast ramp up in HVDC power. This control feature is less commonly applied and it is significantly more challenging than runbacks. The speed and performance of a ramp-up scheme depends on the strength of the AC system. That is, the AC grid must be able to support higher power flow from the HVDC terminal without significant depression of the AC bus voltage. For VSC systems, the HVDC converters can contribute reactive power support within the ratings of the converters. For LCC systems, increased HVDC power results in increased reactive power consumption by the converters, which can cause significant depression of the AC voltage unless additional reactive power can be supplied from other devices (e.g. SVC).

Fast ramp up with short-term HVDC overload. This scheme would temporarily increase the HVDC power transfer above the normal system rating and then return to rated power transfer after a few seconds. This type of scheme could be helpful in arresting the frequency decline during an underfrequency event, during the critical period when generator governors are beginning to respond but have not yet achieved their full response.

Normally, the maximum power transfer of a HVDC system depends on the ratings and capabilities of the various components within the system. For example, converter transformers can withstand significant short term overloads due to their large thermal mass. Converter valves, on the other hand, are constrained by the maximum allowable junction temperature within the power-electronic devices (e.g., thyristors). These devices have very low thermal inertia due to their small size, and junction temperatures increase very rapidly (within a few seconds) when current loading is significantly increased. Such temperature changes are too fast and too localized to be influenced by the valve cooling system.

There is very little experience with short-term (several seconds) overload of HVDC systems. However, one recent example is the Champa-Kurushetra HVDC system[109] in India. The valve design included the following overload capabilities; 31% above rated for 2 hours, 65% above rated for 5 seconds. Of interest, here is the 5-second overload capability which could be beneficial for a FFR function.

All converter valves have some inherent level of overload capability. But if such capability was not specifically required in the specification of the system, the valve's capability is probably unknown. However, it should be possible for the manufacturer to review the design of an existing valve to determine the potential for short-term overload operation. However, this would only be the first step. If overload capability is determined to be feasible, the manufacturer and owner would need to work together to develop a workable scheme to enable the short-term overload functionality while maintaining the reliability and security of the HVDC valve equipment.

Another approach for adding overload capability to an existing HVDC system would be to add another valve group in parallel to the existing valves, thereby increasing the total current rating of the HVDC terminal. The rating of the parallel valves would be coordinated with the short-term overload capability of the HVDC cable. The Pacific HVDC Intertie in the USA was originally built with 2000 ampere valves, and later augmented with 1100 ampere valves in parallel. This approach would be very expensive, and other alternatives such as battery storage systems would likely be more cost-effective.

Speed of Response

Several factors affect the speed of response for rapid changes in HVDC power, especially for increases in HVDC power.

Inherent time constants in the HVDC system itself. Control systems can respond quickly, but the rate of change of current can be limited by the electrical characteristics of long HVDC cables. In a LCC-type HVDC system like BassLink, a severe AC fault in the vicinity of a converter terminal can cause the DC line or cable to discharge, and the DC-side voltage may approach zero. When the AC voltage recovers, the voltage on the DC line or cable must also be restored to near nominal magnitude before the HVDC system can transmit significant power. For the HVDC system to recover, the converters must first inject current into the line/cable to "charge up" the line/cable capacitance and restore DC-side voltage. After that occurs, the inverter terminal can begin injecting current to the ac system, thereby restoring HVDC system power transfer. The time required to charge up the line/cable capacitance, restore voltage, and restore power transfer depends on the system-specific electrical characteristics of the HVDC and AC systems. Commissioning tests of BassLink demonstrated a total recovery response

---

[109] Thyristor valve for the 12-pulse converter for the Champa-Kurukshetra HVDC transmission scheme, Mohammad Hassan Jodeyri; Andrzej Dzus, 2013 IEEE Innovative Smart Grid Technologies-Asia (ISGT Asia), Year: 2013, Pages: 1 - 6, DOI: 10.1109/ISGT-Asia.2013.6698790

time near 300 ms for a DC-side fault.[110] See Figure 41, which is an excerpt from the referenced publication. Although this plot is for a DC-side fault, an AC fault that discharges the DC-side line/cable would exhibit a similar recovery response. And although this response is slow in comparison to other HVDC projects, it is potentially fast enough for a frequency-arresting function.

AC system voltage support. HVDC power transfer depends on the ability of the ac system to maintain the bus voltage at the converter terminal nearly constant. If HVDC power transfer is increased at a rate that exceeds the capability of the AC system voltage support, the voltage will sag and system security will be at risk. System studies are needed to quantify the ability of the ac system to support rapid increases in HVDC power.

Type of control scheme. Fast changes in HVDC power can be initiated by either an open-loop triggered action scheme (e.g., detect breaker opening and initiate predetermined action) or measured-frequency input scheme (a frequency governor or modulation function). (See additional discussion below.)

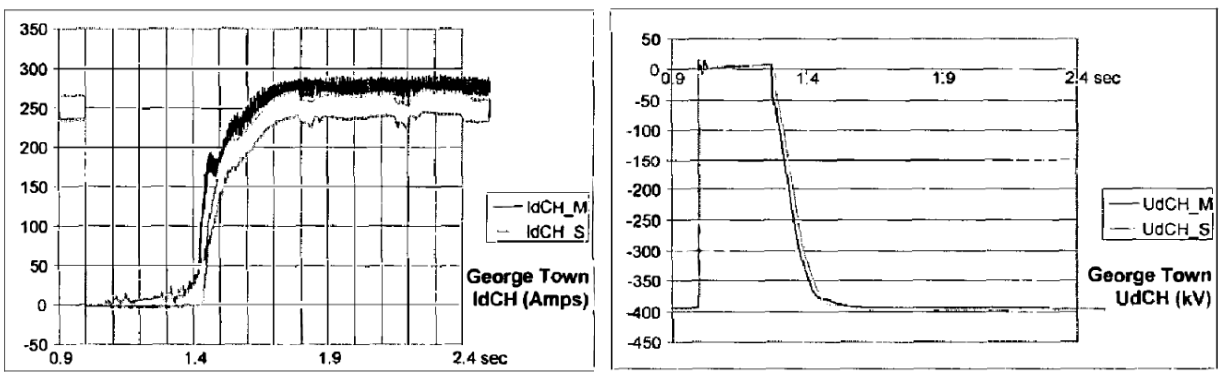

Figure 7: DC Current – High Voltage (IdCH)   Figure 8: DC Voltage – High Voltage Bus (UdCH)

Source: Latest control and protection innovations applied to the Basslink HVDC interconnector, M. Davies et al, 2006

Figure 41 Response of BassLink to DC-Side Fault

3.8.2 Controls

There are several options for implementing frequency arresting controls on HVDC systems:

1. Frequency-sensitive power modulation control function that changes the power order of the HVDC system as a function of grid frequency. This system typically includes a frequency transducer (with voltage and/or current inputs), deadband, gain/phase compensation, and output limits.
2. Triggered action run-back or run-up scheme that changes the power order of the HVDC system by a preset amount in response to a given event (e.g., trip of a specific

[110] Latest control and protection innovations applied to the Basslink HVDC interconnector, M. Davies; A. Kolz; M. Kuhn; D. Monkhouse; J. Strauss, AC and DC Power Transmission, 2006. ACDC 2006. The 8th IEE International Conference on, Year: 2006, Pages: 30 - 35, DOI: 10.1049/cp:20060007

generator or transmission line). Such systems often have multiple triggering events with different power order responses.

3. Combination of both triggered action and frequency sensitive power modulation.

The frequency modulation approach has the benefit of being continuously active, and the magnitude of the response is driven by the magnitude of the disturbance and the grid's frequency excursion. To achieve fast response, these types of controls require high gain and short time constants, which can unfortunately cause undamped oscillations or instabilities under some system operating conditions. The Pacific HVDC Intertie[111] and the Intermountain HVDC system[112] both experienced unstable oscillatory grid events due to such controls when gains were set too high with the intent of achieving fast response. Figure 42 shows one such event in the western USA grid involving the Pacific HVDC Intertie.

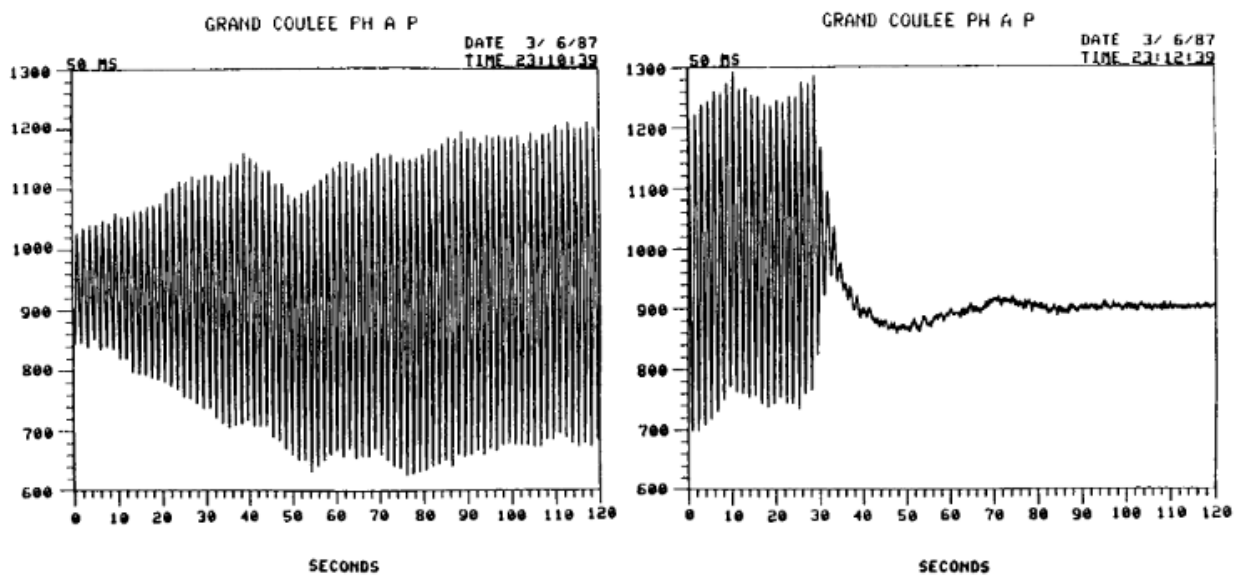

Fig. 1. Participation of Grand Coulee Powerhouse in western system oscillations of March 3, 1987.

Source: Robust damping controls for large power systems, J. F. Hauer, IEEE Control Systems Magazine, 1989

Figure 42 Unstable Grid Event Involving HVDC Frequency-Input Modulation Controller

The triggered action approach initiates an open-loop pre-programmed step change in HVDC power that is pre-calculated to rapidly adjust the grid power balance to more stable condition. It has the benefit of speed (can change HVDC power in a fraction of a second) without the risk of control-loop instabilities inherent with frequency modulation controls. Its response is limited to pre-determined actions for specific triggering events and specific operating

[111] Robust damping controls for large power systems, J. F. Hauer, IEEE Control Systems Magazine, Year: 1989, Volume: 9, Issue: 1, Pages: 12 - 18, DOI: 10.1109/37.16744.

[112] 1994 System Disturbances -- Review of Selected System Disturbances in North America, North American Electric Reliability Council, October 1995.

conditions (e.g. flow on an AC interconnector). Individual triggers can be armed when needed and disarmed when not needed.

The combination approach uses a triggered action to rapidly change HVDC power to an appropriate post-disturbance level with a slower-acting frequency modulation control to fine tune the post-disturbance operating point. This is a promising approach when high-risk grid events are well known and directly observable (i.e., trip of a line or generator).

New Zealand, Pacific HVDC Intertie, Sardinia-Corsica-Italy HVDC, Intermountain HVDC, Gotland HVDC, Square Butte HVDC, and many other systems have power modulation controllers that respond to changes in grid frequency. In parallel AC/DC systems, they provide damping to inter-area oscillations and in some cases, improvement in first-swing stability following disturbances. For isolated or island systems, the HVDC provides governing response like a generating plant – i.e., contribute to frequency regulation using droop control. In these systems, the controllers were custom-designed using dynamic simulation tools to design system specific settings to balance dynamic response and stability over a wide range of system operating conditions.

3.8.3 Possibilities and Trade-Offs with HVDC for Fast Frequency Response

The preceding discussion related to the physics and experience with HVDC systems performing functions that could contribute to fast frequency response. The following discussion is intended to provide some overall perspective with respect to the challenges faced by the AEMO grid.

The most critical events leading to rapid frequency decline is for loss of the Heywood Interconnector) which leaves South Australia as an electrical island with a generation shortage. This event is most severe when flow on the Heywood interconnector is at or near maximum power flow from Victoria to South Australia before the event. When Heywood is heavily loaded, MurrayLink is likely to be loaded at or near its maximum as well. This limits the possibilities for MurrayLink loading to be transiently increased to help mitigate the South Australia unbalance.

If frequency decline is rapid, then the response time needs to be as fast as possible. The total response time would include the time to detect the event, the time to communicate that information to the HVDC terminal, and the time for the HVDC system to respond.

If the event is not initiated by an ac fault (e.g., if the ac tie simply trips or if a South Australian generator trips due to an auxiliary system problem), then the response time for a step change in HVDC power could be very fast, perhaps within 50 to 100 ms. However, if the event is initiated by an AC system fault, then there is the possibility that the HVDC system voltage might be affected and that the DC-side line voltage would need to recover before HVDC power can be moved to a new (higher) level. HVDC system response time would be much longer that 100 ms (perhaps near 300 ms).

Given that the HVDC response (due to physics) may be slower than desired, it is therefore critical to reduce the event detection and communication times as much as possible in order to minimize the total response time. This is commonly achieved in remedial action schemes using the following approach:

- Detection: Monitor the opening of critical breakers using "B" contacts on the breakers themselves, or breaker status information in protective relays. Detection time is less than one cycle from breaker opening.
- Communication: Send breaker status signal by dedicated communication channel to the HVDC terminal, using communication systems that interconnect the transmission substations. Similar remote status signals are commonly implemented as part of transmission protection schemes, and have used power-line carrier, microwave, or fiberoptic communication channels.

Some systems have implemented frequency-input governor type schemes with the intent of eliminating the need for communications. However, frequency input systems introduce additional challenges that affect the total speed of response.

Frequency transducers use AC bus voltages as inputs, and the voltage waveforms become very distorted during AC system faults and during the recovery period after the fault is cleared. To keep the transducer from producing erroneous outputs, it is necessary to slow down or filter the transducer output signal.

If fast response is needed from a frequency input controller, then the controller gain must be high and the time constant must be short. But excessively high gains and/or short time constants reduce the closed-loop stability margins to a point where widespread spontaneous or growing oscillations may occur on the power grid. To avoid such oscillations, the gain must be limited, which also limits the speed of response.

The bottom line is that there is no obvious "best choice". A reasonable path forward is to gather the necessary information to quantify each of the factors discussed above and to then select the most promising approach. The first step would be to initiate a study that models the performance of the MurrayLink HVDC system (using EMTP or PSCAD type models), quantifies their response characteristics, and determines their capability to perform additional functions to assist with fast frequency control.

How fast can it respond? Speed of response depends on two factors that are unique to each application.

- The design of the HVDC system. Back-to-back HVDC systems with strong AC systems can respond in 50 ms. Long HVDC cables connected to weak AC systems have response times in the range of 200 ms to 500 ms.
- The type of initiating event. If the grid event does not involve an AC fault near the HVDC terminal, HVDC response time is very fast, in the range of 50 to 100 ms. If the grid event includes an AC system fault that discharges the HVDC line or cable, then the

response time includes the fault recovery time of the HVDC system, which is in the range of 100 ms to 500 ms, depending on the strength of the AC system.

How much can it inject or absorb? Any power level within the rating of the HVDC converter equipment, including any short-time overload capability.

How long can it sustain this response? Indefinitely, for any power level within the HVDC converter ratings. Short term overloads are limited by the thermal time constants of the semiconductor devices, which are typically in the range of 2-5 seconds.

Is there a particular ramp or shape to the response? No. But the rate of response may need to be constrained if the AC system has insufficient dynamic voltage support to prevent collapse of the ac bus voltage when the HVDC power is rapidly increased.

What determines each of these limits? Is there potential for it to be adapted? Lots of factors determine the limits. Response is highly system-dependent per the preceding discussion. Detailed EMTP-type simulation studies are required to determine actual performance capabilities for the MurrayLink HVDC system.

Where, when and how have these capabilities been demonstrated? HVDC power runback schemes are very common and are offered as standard feature on most HVDC systems. New Zealand, Pacific HVDC Intertie, Sardinia-Corsica-Italy HVDC, Intermountain HVDC, Gotland HVDC, Square Butte HVDC, and many other systems have power modulation controllers that respond to changes in locally measured grid frequency. The authors are not aware of any specific systems that include HVDC power run-up (step increase) functions, but the implementation would be the same as for runback. However, schemes to rapidly increase power require more technical analysis to quantify performance capabilities and limitations.

What is the impact on the resource from responding in this manner? Provide a quantitative estimate of any immediate (short run marginal) cost. No impact on HVDC converter equipment, provided the control actions are designed to keep currents and voltages within the rated capabilities of the equipment.

How mature is this technology? Is there potential for these characteristics to improve in the future? To what degree, and over what timeframe? Both triggered response functions (runback/run-up) and frequency modulation functions are mature technology. Use of HVDC short-term (2 to 5-sec) overload is less mature, as very few systems have included such capabilities in their design specifications. The Champa-Kurushetra HVDC system is the only example found during the search of published literature.

How is the ability to deliver a fast frequency response service affected by nearby power system disturbances, such as transmission faults? Will the FFR service still be available in the period following a transmission fault? How could this be managed? Nearby AC transmission system faults have significant impact on:

- The speed of response of the HVDC terminal, since the HVDC must recover to prefault power flow before it can provide additional power for FFR.

- The ability to rapidly and accurately transduce frequency after the fault, due to voltage distortion caused by the ac fault and by the subsequent recovery of the HVDC converters.

In summary, the preponderance of industry experience suggests that the most effective and secure approach for HVDC to contribute to FFR is a remedial action scheme (RAS) that detects the most critical events via breaker status, transmits a triggering signal to the HVDC terminal, and initiates a pre-programmed action designed for the specific event. With this approach, event detection and signal transmission happens in parallel with the ac and HVDC system fault recovery, and by the time the HVDC converter is recovered and able to respond, the communicated signal arrives and triggers a specific open-loop action. This scheme enables a large response in a very short time, and avoids the significant risk of widespread oscillatory instabilities associated with frequency-input HVDC power modulation schemes.

4 IMPLEMENTATION RISKS AND CONSIDERATIONS

This section discusses a range of practical considerations which may be common to a solution of the frequency control problem.

4.1 Torsional Impacts on Turbine-Generators

Very fast active power injection by any inverter-based FFR resource will cause torsional impacts on synchronous machines operating in the electrical vicinity in the power system. Amplitude and rate of rise of power injection, coupled with electrical proximity to operating synchronous machines dictate the degree of impact. Due diligence requires that potential risks, including potential loss of life, especially for steam units, be considered. These concerns apply, generally to a lesser extent, to wind generation, particularly type-1 and type-2 machines. Transient torques from open-loop step inputs, and closed-loop stimulation of torsional oscillatory modes are both risks. Both of these concerns are particularly acute if there is risk of frequent inadvertent triggering of response.

For systems using frequency-input power modulation, there is a risk that the controller may reduce the inherent damping of torsional vibrations to a point where spontaneous growing torsional oscillations[113] may occur. [This was the case for the Square Butte HVDC transmission system in the USA[114]. It was necessary to modify the control parameters and include torsional notch filters to reduce the interaction to safe levels]. The design of frequency-input modulation controllers should investigate the risk of torsional interaction in advance, and if torsional interaction risk is found to exist, mitigation and protection features must be included in the design.

4.2 Fault induced voltage and power depression

Grid events involving large frequency excursions are often triggered by faults in the AC transmission system. Aside from the temporary interruption of AC power flow, the faults produce a few other detrimental consequences that affect FFR:

- Operation of power electronic devices in the vicinity may be seriously disrupted. HVDC converters as well as converters for FFR devices may experience power flow disruptions to low or distorted voltages, both during the fault and for some period after the fault as the converter control systems resynchronize to the post-fault system voltage.
- Transducers measuring frequency or RoCoF are more likely be unable to accurately determine grid frequency during faults and for a short period after fault clearing due

113 The turbine-generator drive-train of synchronous thermal generation and wind turbines have multiple masses – at least the turbine and the generator - connected by a shaft. Oscillations of this mass-spring system put stress on the shaft.

114 M. Bahrman, E.V. Larsen, R.J. Piwko, H.S. Patel, "Experience with HVDC - Turbine-Generator Torsional Interaction at Square Butte", IEEE Transactions, PAS, Vol. 99, May-June 1980, pp. 966-975.

to the distorted voltage waveforms that exist during the fault and during the subsequent recovery of the power electronic converters.

Quantifying these effects would probably require a system study using EMTP-type models to simulate the fault events and determine the voltage waveforms that occur during and after the fault events. To do this right would likely require a massive EMTP study of the South Australia region.

As noted in Section 3.8, the vulnerability of the HVDC to the voltage depression associated with a fault that trips the AC tie line puts the power from the DC line at risk just as the system needs it most. That is, HVDC power may be interrupted or reduced during the fault and it would need to recover before providing any contribution to FFR. Of the various technologies discussed in Chapter 3, the large (non-inverter type) flywheel technologies have particularly good characteristics for reducing this risk. Colocation of these devices in the electrical proximity of the HVDC terminals could have benefits for the fault recovery performance of the DC.

Wind turbines, PV and other devices may well have blocked power flow during significant voltage depressions. Physics limits the amount of power that can be injected into a low voltage. That makes the overall FFR issue more complex as well. For all the subject technologies, guaranteeing fast response in light of this complex behavior adds to the challenge. Further, the ability of the various inverter-based FFR technologies to start (or resume) power injection immediately after a fault will vary with technology and equipment design. In general, as mentioned above, devices that are more resistant to change in the DC link voltage of the inverter, such as batteries, will tend to be able to inject power sooner following AC voltage recovery.

For additional context, EirGrid and SONI are implementing a new system service product to quickly recover MW output following a disturbance: Fast Post-Fault Active Power Recovery service. Following any fault that is cleared within 900 ms, including transmission faults, a plant providing this service must recover its active power to at least 90% of its pre-fault MW output within 250 ms of when the voltage recovers to at least 90% of its pre-fault level[115]. This service is a market mechanism to encourage faster recovery than the minimum required by the grid code. It was created in recognition of the difficulty of meeting this requirement, and is also an incentive for existing (e.g. grandfathered) resources to invest in speeding up their performance.

EirGrid/SONI's phase 2 RoCoF study finds that for Ireland to maintain RoCoF within 0.5 Hz/s, they need synthetic inertia to begin responding by 100ms from the start of the event (as visible to the device) and to achieve full response within 200 ms from when the response begins[116]. This has not been translated into a service yet.

115 EirGrid and SONI, *DS3 System Services: Portfolio Capability* Analysis, November 2014.

116 EirGrid and SONI, *RoCoF Alternative & Complementary Solutions Project Phase 2 Study Report*, March 31, 2016.

4.3 Voltage control and collapse (weak AC system)

Rapidly increasing the power output from any system device may impose a risk if the system is not adequately "strong" or if it does not have adequate fast-acting voltage support to maintain the bus voltage. Increasing power flow without adequate voltage support can lead to local voltage collapse, and in severe situations, cascading voltage collapse.

This is a problem that is explicitly considered in the active power behavior and recovery of wind plants and HVDC systems. Normally in weak grids, the active power recovery is *deliberately* kept slower than the reactive power recovery (per the preceding paragraph) for the reason that the power often cannot be evacuated from the wind plant or HVDC terminal to the loads until the voltage has been restored. Thus, in weak systems, the active power recovery tends to be slower. This is fundamentally at odds with the challenge facing South Australia, where fast power recovery is necessary for FFR.

An assessment of the power recovery performance of FFR devices under consideration would need to be undertaken by EMTP-type simulation analysis of the South Australia grid. This "weak system" performance analysis could be performed as part of the post-fault recovery analysis discussed above.

4.4 Differences between Under-Frequency and Over-Frequency Response

The ability of a generator to provide an over-frequency response is different from under-frequency response. Over- and under-frequency response entail different costs, can have different speeds, and can provide different levels of aggressiveness. As such, these two responses should be disaggregated and considered two different products or services, similar to up-regulation and down-regulation in some ISO's.

Figure 43 shows the droop characteristic for a plant. Consider a plant operating at the intersection of the dashed blue (nominal frequency) and dashed green (dispatched set point) lines. As frequency increases above the deadband, the slope of the droop characteristic tells the plant how to decrease output. It is noted that if the plant is operating at a minimum generation level, then the plant may not be able to decrease output unless it is decommitted. As frequency decreases below the deadband, the slope of the under-frequency droop characteristic tells the plant how to increase output. Note that the plant needs to be operating below its maximum output in order to have headroom to increase output.

Over-frequency response has no opportunity cost and little or no variable cost. It is beneficial to the system in the event of a loss of load event, during some system separation events, or in response to overly aggressive UFLS on a loss-of-generation event. This last consideration is not insignificant, as the discussion of backswing of frequency illustrates. To the extent that the cost to enable fast over-frequency response is very small, enabling this response at all times should be considered.

The aggressiveness of the response (droop) may vary for over-frequencies versus under-frequencies. In the past, this response was symmetric. This was largely a historical artifact of mechanical (e.g. flyball) governors that were symmetrical (the weight was fixed). But modern controls, and particularly those available on wind or PV plants, need not be symmetrical. Droop characteristics may be different for over versus under-frequency for wind, and especially PV (since there are no mechanical stress limits on how fast PV can back down). The system should take advantage of these responses.

While the primary focus of PFR tends to be on under-frequency events, as was noted above, historically over-frequencies can lead to, or be part of, severe system disturbances (e.g. Florida's Turkey Point blackout in the USA, Malaysia blackout). Unlike combustion-based resources (especially gas turbines), wind and, especially solar PV, can reduce power output fast and to very low levels without much risk of tripping offline. Such aggressive response requires higher gains (less droop) for over-frequency. The recommended approach is to explicitly *allow* (not require) that generating resources can have asymmetric droops and that they can be less than 5%, based on mutual agreement of the grid operator and the plant owner.

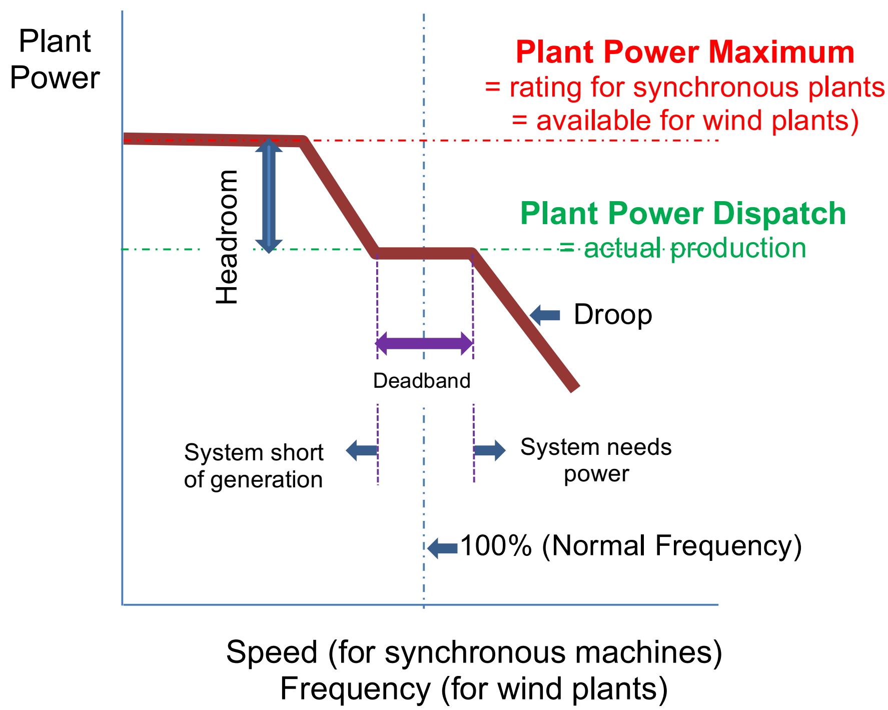

Figure 43 PFR droop characteristic.

4.4.1 Risk of Conflicts with PFR Droop Response

It is important to remember that this figure is static, and the response is linear. It is possible to create or tune the dynamic response of generation such that the ability of the unit to reasonably follow this droop characteristic is degraded or defeated. For example, functionality that focuses on achieving specific power ramp rates or rapid response to market signals can conflict with droop control. Ill tuned controls can result in frequency correction overshoots, or in sustained frequency errors. For example, in the US, generator load controls that defeat ("withdraw") PFR are presently in widespread use. These cause poor frequency response, especially in the US Eastern Interconnection. Efforts are currently being made to correct this problem (by NERC). For any system to evaluate FFR properly, dynamic models of PFR need to correctly capture all these effects.

4.5 Implications and Practicalities of Zero Synchronous Inertia Operation

While the focus of this document is limited to technologies that will limit RoCoF and hold frequency, the trends of South Australia mirror those of other places on a trajectory towards minimal and even zero synchronous generation. This is the subject of intense interest in research circles, as evidenced, for example, by a session at the July 2016 IEEE Power & Energy Society General Meeting dedicated to the topic. Broadly stated, the trends in development are towards inverter-based resources that tend to mimic synchronous machines, in that they synthesize an internal voltage phasor that sets both magnitude, angle and frequency. The example shown in Section 3.3, Figure 24, illustrates that 100% inverter based operation of an island is possible. In that case, a single device, specifically designed for this purpose, provided all the needed functionality for the island. The on-going developments are centered around how to get devices to share duty, especially current; how to get coordinated frequency setting; and how to deal with greatly reduced fault currents and different fault current behaviors for protection. Rapid development over the next few years is anticipated.

However, it is important to note that there are good reasons for the present design and control of the voltage-sourced current controlled inverters that are predominant today for both wind and solar PV generation. These devices make economically effective use of the relatively expensive power electronic components; they tend to naturally coordinate well (between many devices); they tend to have stable and relatively benign interaction with other system components, including generally better damping of system oscillations. For the inverter-based power generation industry to make wholesale movement towards equipment that more closely mimics synchronous generation has non-trivial cost and systemic performance issues that need to be addressed. This is not just a simple question of adjusting control algorithms.

4.5.1 WECC and other Research on 100% Inverter-based Systems

There is good research being done in this arena. Of particular note, a new paper[117] looking at a hypothetical future US Western Interconnection with 100% inverter based generation shows great promise for this (virtual synchronous machine) approach. Some good work is coming out of National Grid in the UK.

---

117 Transient Stability Analysis of an all Converter Interfaced Generation WECC System; Deepak Ramasubramanian, Vijay Vittal, John M. Undrill. IEEE GPM 2016

A new initiative in the European Union, called Massive InteGRation of power Electronic devices (MITIGATE), includes participation by 13 countries, 12 TSOs and is funded for about 17M Euros. The scale of the project (and funding) gives a sense of the breadth of the challenge.

4.6 European Rooftop Solar Experience

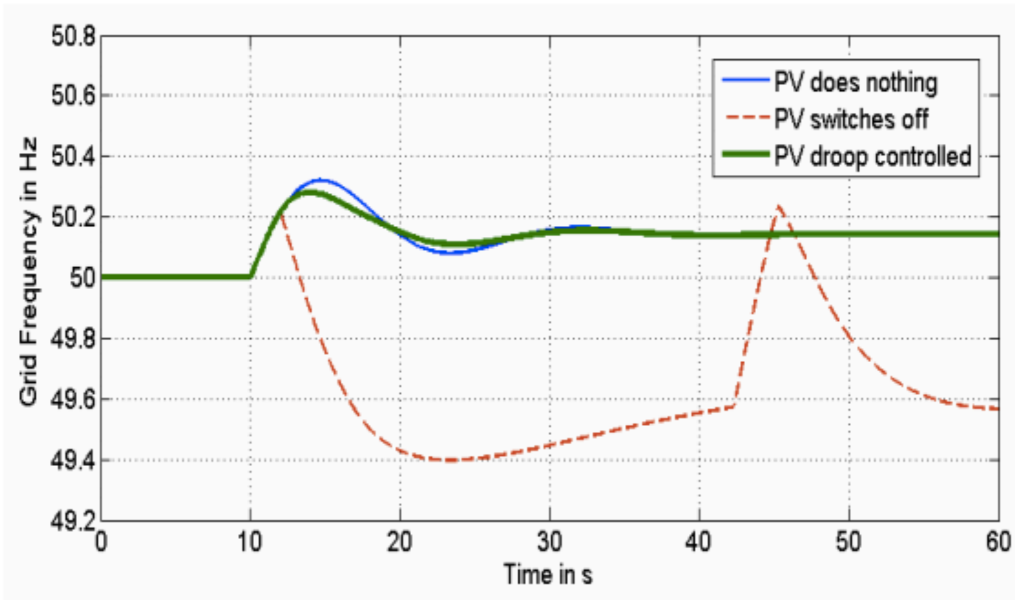

Figure 44 Overspeed event due to loss of load of 5 GW in Europe resulting in disconnection of 10 GW of DER at 50.2 Hz[119]

Germany has extensive experience with retrofitting DER to comply with updated connection requirements. About 12,700 MW of rooftop solar was connected to Germany's low-voltage distribution system. During 2005-2006, it was not anticipated that there would be high penetrations of DER, so the connection requirement was established such that the rooftop solar inverters must trip offline for over-frequency events at 50.2 Hz120. The European grid was designed to survive a sudden loss of load of only 3,000 MW, much smaller than the amount of generation that could be lost upon common-mode tripping of rooftop solar. Unlike voltage, frequency is, for the most part, the same across the interconnected power system. A loss of load (such as loss of an intertie during export conditions) could increase frequency to 50.2 Hz, which could result in much more than 3,000 MW of rooftop solar (depending on solar resource at that time) common-mode tripping. The dashed red trace in Figure 44 shows a simulation of such an event in the European grid, with the rooftop solar systems tripping offline

119 Kastle, G. and Vrana, T., http://www.cired.net/publications/cired2011/part1/papers/CIRED2011_1275_final.pdf

120 J. Boemer, et al, "Overview of German Grid Issues and Retrofit of Photovoltaic Power Plants in Germany for the Prevention of Frequency Stability Problems in Abnormal System Conditions of the ENTSO-E Region Continental Europe," *1st International Workshop on Integration of Solar Power into Power Systems*, Aarhus, Denmark, Oct. 24, 2011.

at 50.2 Hz and then trying to reconnect after 30 seconds. The green trace simulates the same event, but with the PV providing over-frequency droop response.

Nine hundred distribution system operators had to come together to reach agreement to retrofit some 350,000 existing systems at an estimated $520 M cost to correct this reliability problem[121]. Being proactive and undertaking careful planning is one of the lessons learned from the German experience.

There are two issues at play here. One is that the existing connection requirements should be revised so that DER supports system security – so that DER does not common-mode trip in response to a small voltage or frequency deviation. The second is whether DER should be required to provide FFR capability. The new Australian standard AS4777, which came into effect this year, requires frequency response from distributed PV. Presumably, the continuous nature of this new standard will avoid the problems created by the discrete control requirement in Germany. There is hysteresis in AS4777, which should avoid cycling risks. Complex dynamic events, involving both high and low swings of frequency may create unexpected behaviors. But whether DER should be capable of under-frequency response is a more open question. On the one hand, DER is growing quickly, and as Germany has seen, retrofits can be expensive, time-consuming, and resource-consuming. It is unlikely that rooftop PV owners will be interested in pre-curtailing their systems to try to make money in a potential FFR (or PFR) market. However, if one believes that very high DER penetrations or micro-grids are realistic future scenarios, then proactive planning would suggest that requiring FFR in the connection rules for future small resources might be prudent. But great care must be taken to avoid unintended consequences.

[121] B. Ernst, SMA, “Evolution of LV PV Interconnection Requirements in Germany,” Utility Variable Generation Integration Group Spring Technical Workshop, Anchorage, AK, May 2014.

5 INTRODUCTION: FFR REQUIREMENTS AND ANCILLARY SERVICE SUITABILITY

The situation in South Australia, while extreme, is not unique. The decline in synchronous inertia and displacement of traditional sources of frequency sensitive generation that is accompanying the addition of large amounts of wind and solar generation to bulk power systems is the subject of considerable concern in the industry. Increasingly, entire grids are faced with the prospect of operating under conditions where the majority of system generation is from inverter-based, i.e. wind and solar, generation. Additional resources, such as energy storage and HVDC, are growing. As the traditional sources of primary frequency response (PFR), i.e. governors on the turbines of synchronous generators become scarcer under some operating conditions, new sources of frequency response are emerging. Much has been written about inertia-based controls for wind turbines, fast acting load controls, battery energy storage and a diverse host of other technologies and concepts.

For system planners and market designers, evaluation of the system frequency for these conditions is new ground. As described in 2.1, it is well understood that the initial rate of change of frequency (RoCoF) increases (becomes more negative) with dropping synchronous inertia. Indeed, the calculation of initial RoCoF is one of the few quantitative relationships that is relatively simple to provide in closed form. But beyond this, the relationship between traditional PFR resources and new resources is complex. In addition to PFR, the resources, described in depth above can provide FFR, which can be more effective than a similar quantity of PFR. Again, in general, it is well understood that with declining synchronous inertia the time available for frequency responsive resources to deliver the arresting power necessary to create an acceptable frequency nadir declines. However, the diversity of behavior from the emerging resources presents a quantitative challenge. The planner and market designer must "have enough" of these resources available to meet the system needs.

Several systems, notably ERCOT and EirGrid, have made inroads into understanding the problem of "how much is enough?" Engineers in those, and other, organizations have spawned the introduction of a variety of fast frequency response ancillary services. In the US, ERCOT has lead development of understanding of the relationships between system inertia, PFR and FFR. In this section, we use a very simplified model of the South Australia system to *illustrate* the analysis necessary to answer questions like "How much is needed?" and "What is most effective?"

5.1 Arresting Power and Energy

For the majority of power systems, including the major interconnections of North America, frequency response requirements in the first seconds following major system disturbances drive primarily towards one overriding objective: avoiding involuntary UFLS.

To that end, great attention is being given to understanding and assuring that sufficient arresting power is provided to stop frequency decline before specific levels of under-frequency are broached. The fundamental concepts and needs have been well documented.

The NERC Frequency Response Initiative Report[122] provides arguably the definitive (at least for North America) discussion. A number of studies by the authors[123], [124], [125], as well as other work in the industry have shown that a variety of frequency responsive technology beyond traditional reliance on governor response from synchronous generation can be highly effective in providing this essential arresting power.

5.1.1 Frequency Nadir: the metric of primary frequency response

There are a variety of complexities, but overall a single measure of performance captures how well a system is performing: the depth of the frequency nadir following a disturbance. For most systems there are specific credible ("design-basis") disturbances for which involuntary UFLS must be avoided. Various other measures are used to help determine how well the system is meeting that objective. For example, settling frequency and frequency response are, respectively, a proxy for how well the frequency was stabilized and an indicator of how various parts of the system contributed to the outcome. One measure that gets lots of attention is RoCoF. RoCoF, for most systems and major interconnections is, in itself not a performance objective, but rather a measure that is useful to help understand what behavior (especially speed) is needed.

The focus here is exclusively on the ability of resources to improve frequency nadir following loss of generator or infeed disturbances. Most of the concepts presented should extend to RoCoF objectives, in the rare system where managing RoCoF is itself a legitimate performance objective. (Discussion of the vulnerability of power system components to high levels of RoCoF is addressed in a separate report[126]).

For discussion of FFR, with particular attention to South Australia, Figure 5 from above is reproduced here in Figure 45, with additional annotations on the FFR. In the figure, the total power signature of a resource providing FFR is broken into four distinct regions. The energy delivered before reaching UFLS is noted in green. The *only* arresting energy that is useful towards avoiding involuntary UFLS must (obviously) be delivered before the system frequency reaches that threshold. Arresting energy delivered after that time, shown in orange, is still useful in managing the depth and timing of the frequency nadir. In systems with staged UFLS, the distinction between the green and orange becomes less distinct. Ultimately, the modelling presented here is based upon the objective to avoid involuntary UFLS, and the results presented here are based on the concept of keeping the frequency nadir above a known threshold. The third area, in blue, represents the period during which the FFR should be being replaced by PFR. This must be done slowly enough that the frequency doesn't drop again.

---

122 NERC (2012c). Frequency Response Initiative Report. Draft. September 30. http://www.nerc.com/docs/pc/FRI%20Report%209-30-12%20Clean.pdf

123 N. Miller, et al, "California ISO (CAISO) Frequency Response Study," Nov. 9, 2011. https://www.caiso.com/Documents/Report-FrequencyResponseStudy.pdf

124 N. Miller, et al, "Eastern Frequency Response Study," NREL//SR-5500-58077, May 2013. ttp://www.nrel.gov/docs/fy13osti/58077.pdf

125 N. Miller, et al, "Western Wind and Solar Integration Study Phase 3 – Frequency Response and Transient Stability", NREL/SR-5D00-62906, December, 2014. http://www.nrel.gov/docs/fy15osti/62906.pdf

126 New RoCoF advisory report by GE. Planned release December 2016.

And finally, for some resources, most notably wind inertia-based FFR, the arresting energy provided must be recovered by reducing the total power delivered to a level below pre-disturbance.

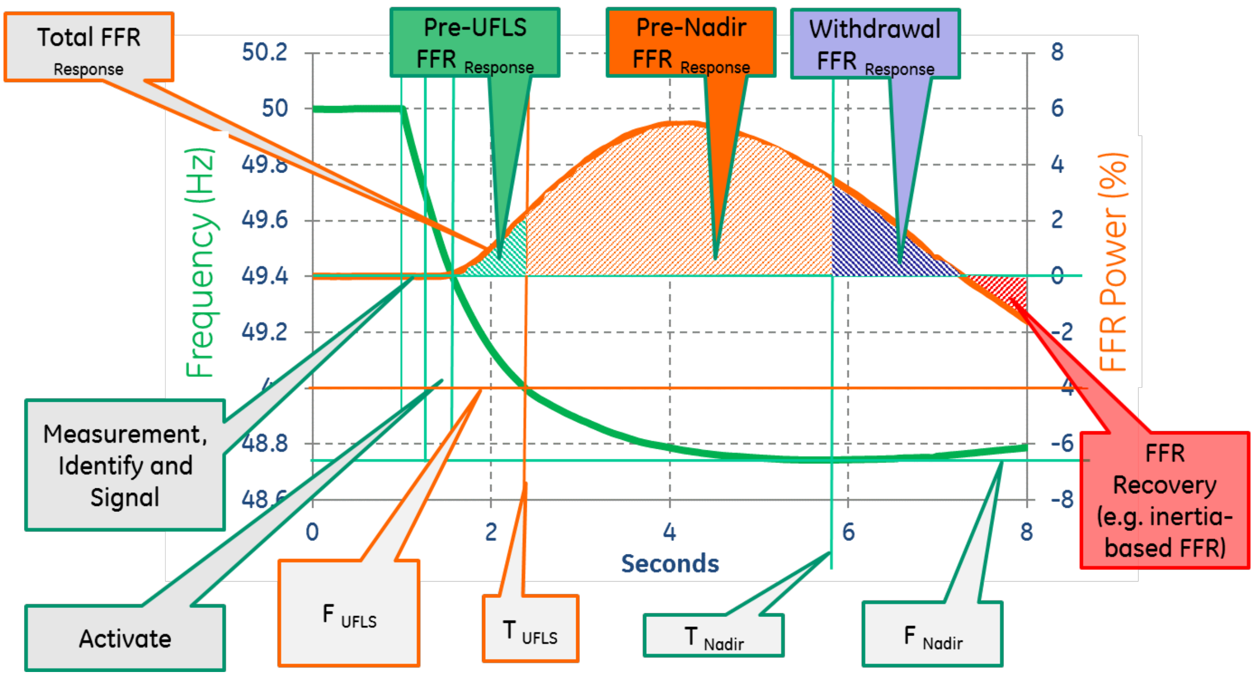

Figure 45 Time Elements of FFR

6 ILLUSTRATION OF SOUTH AUSTRALIA FFR REQUIREMENTS AND ANCILLARY SERVICE SUITABILITY

In this section, we provide an analytical demonstration of the efficacy and requirements for a FFR service in the context of meeting AEMO system performance requirements. The intent of this work is to quantitatively show the relationships between synchronous inertia, conventional primary frequency response services and possible new fast frequency response. The analysis uses a very simple and approximate system representation, and is intended to give results that are quantitatively meaningful, but not precise. The intent here is illuminate processes by which in-depth analysis and more detailed models can be used to plan the system and to guide procurement of ancillary services and operation as it relates to frequency response.

In no way, should the reader regard this illustrative analysis as a substitute for detailed, high fidelity modeling and analysis of the system. Nor should the reader regard this as an exercise in replicating any specific event in South Australia, or validation (or invalidation) of the present UFLS practice. We are attempting to look forward here.

6.1 Discussion of Model and Events

This section describes very simple representation used for the investigation. The model was developed based on inputs from AEMO, and assumptions made by the authors.

6.1.1 South Australia Loadflow Model

A sketch of the system representation is provided in Figure 46. Equivalents of thermal and wind generation, and South Australia system load are indicated. Since primary frequency response is such a critical element to the issue, the thermal plant equivalent is segregated into an equivalent of the units providing primary frequency response (i.e. those which have governors enabled and which will respond to a drop in unit speed), and those that are not responsive (i.e. either have governors inactive or which are at maximum power or valves wide open). The distinction is important, and must be correctly made in actual system studies. The transmission system is largely ignored, except for an explicit equivalent of the Heywood Interconnector, as shown. The MurrayLink power injection is separately represented, with a highly simplistic power injection (i.e. there is no dynamic model of the DC included). As noted, the model is intended to give insight into frequency response.

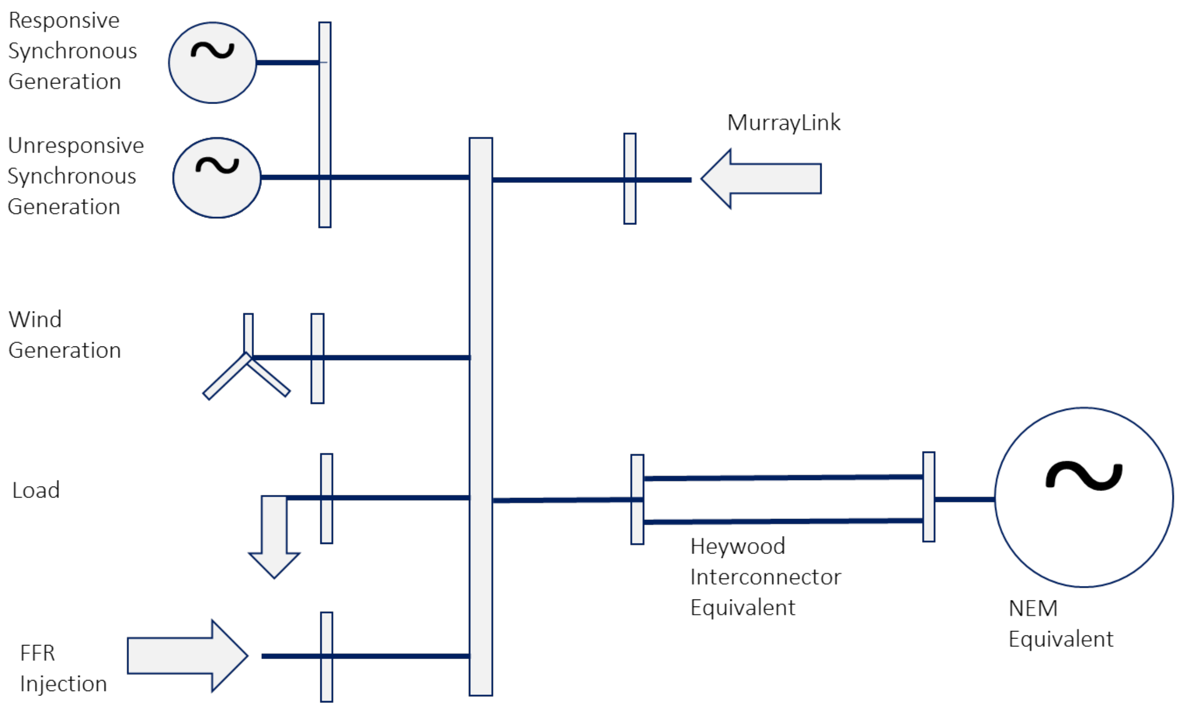

Figure 46 Sketch of Illustrative System Model

6.1.2 South Australia Dynamic Model

The dynamic model of the system is a fundamental frequency positive sequence representation, of the type commonly used for system planning on software platforms like PSS/e or PowerFactory. The cases shown here were run on the PSLF software platform. Representation of the synchronous generation includes standard 6 state machine models and excitation systems. As noted, a governor model is included on the responsive thermal equivalent. The wind turbine equivalent uses a single representation of a typical, present commercial utility scale wind turbine generator (WTG), with no frequency sensitive controls enabled (unless otherwise noted below). The load model includes voltage sensitivity and a dynamic component based on standard load modeling used in the US Western Interconnection planning studies. The load is represented in granular pieces, to allow for approximation to the UFLS system in place in South Australia.

6.1.3 November 1, 2015 Event

To calibrate the simple model, it was tuned to give a plausible response to the November 1, 2015 separation event. The loadflow was adjusted to approximate the grid condition before the event, as shown in Figure 47. The system inertia in South Australia is modeled at 9000 MW-sec for the event. The AC import for the case, and therefore the size of the event for simulation of the event, is 232 MW on the Heywood Interconnector.

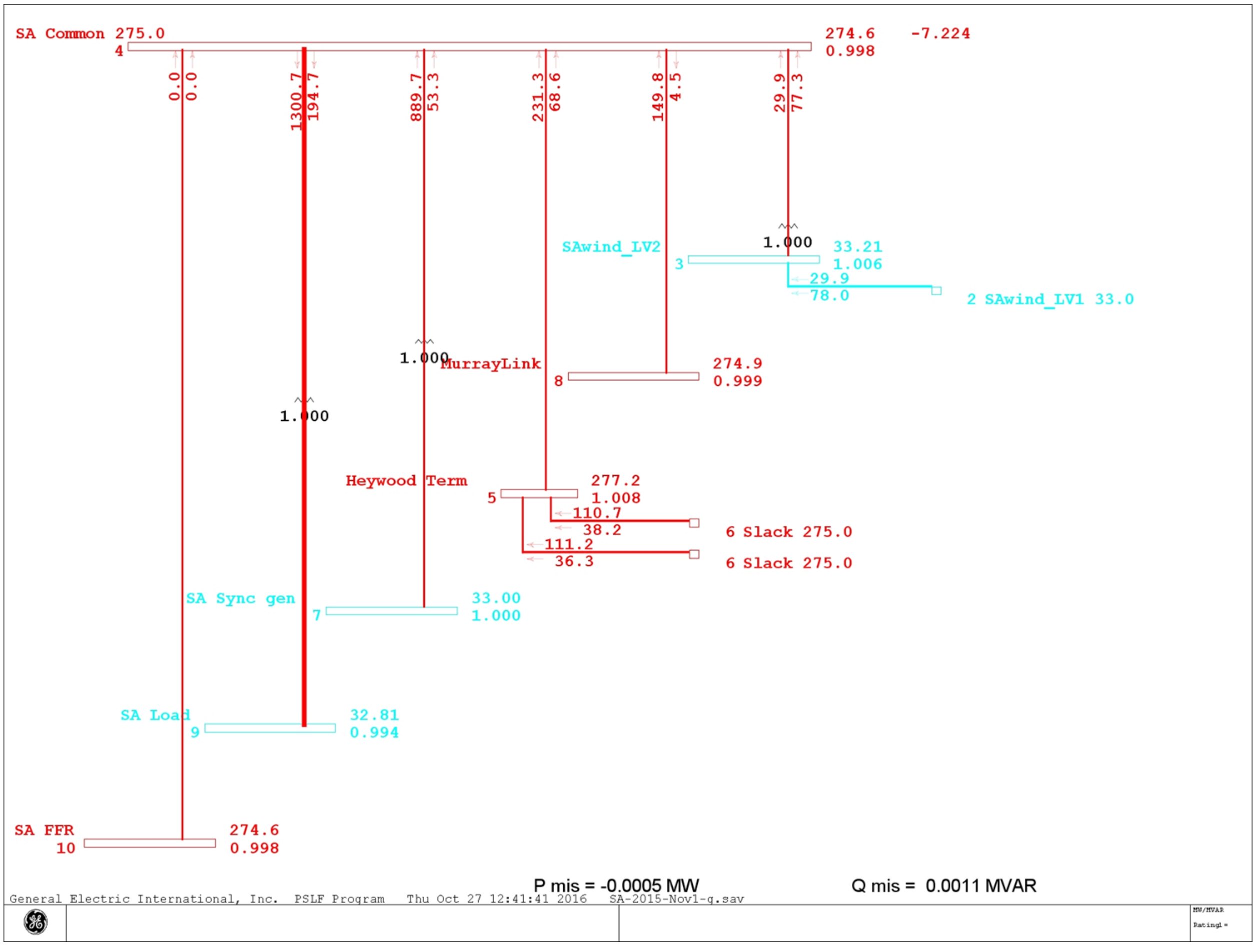

Figure 47 Loadflow for November 1, 2015 event cases

The turbine governor and underfrequency load shedding was tuned to give a reasonable approximation to the measured event. A comparison of the simulated response and a measurement of the event is shown in Figure 48.

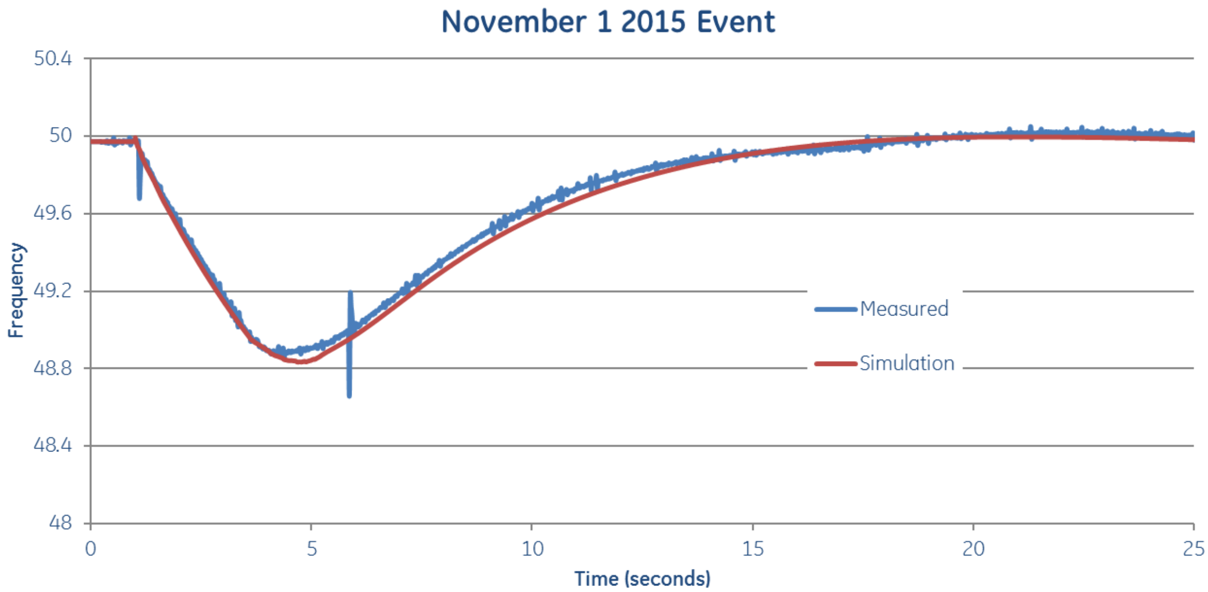

Figure 48 Measurement and Simulation of November 1, 2015 Event

The match is quite good, and provides a basis on which to do the illustrative work. The reader is cautioned to remember that tuning a model to a single case does not constitute validation of the model, which is considerably more difficult.

6.2 Requirements for South Australia: Present Low Inertia System

This section examines requirements for frequency response in the South Australian system as it is operated most of the time today.

6.2.1 Recent History for South Australia

AEMO has shown that there has been a general trend of dropping synchronous inertia in the South Australia system over recent years. The trend has continued in 2016. Figure 49 shows data provided by AEMO up through September 26 of this year, presented as a duration curve. The curve shows inertia down to a bit below 3000 MW-sec.

In this section, performance and behavior down to this level will be examined. Later, the investigation will look at even lower levels.

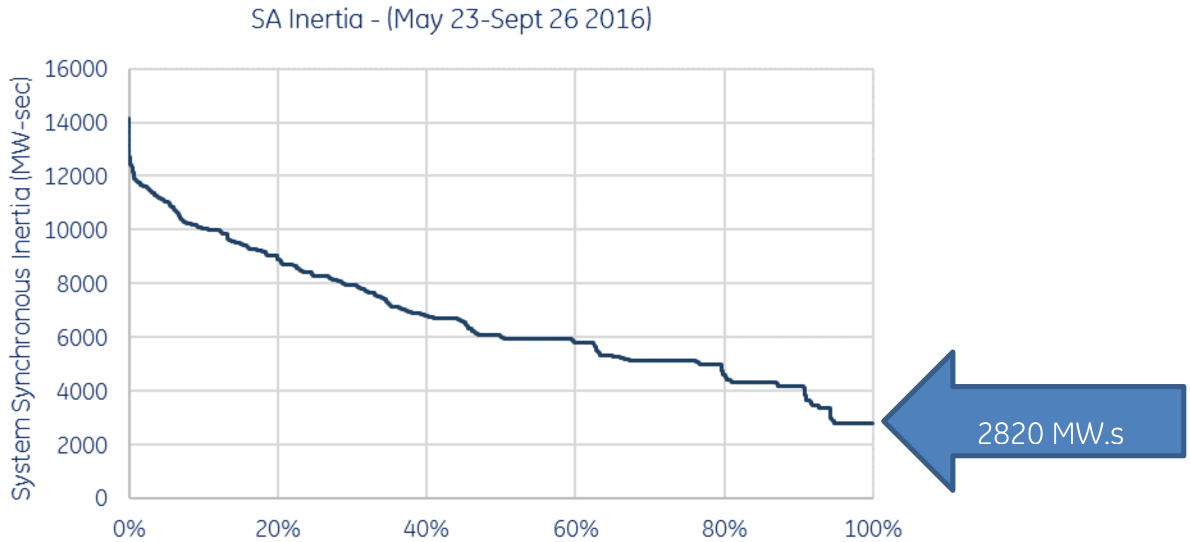

Figure 49 Recent (2016) History of Inertia in South Australia

The severity of frequency events is dominated by the amplitude of the generation-load imbalance. For this discussion, that corresponds to the amount of power being imported to South Australia on the Heywood interconnector. A duration curve of the AC imports (this does not include the flow on the DC MurrayLink) is shown in Figure 50. During the period reported here, the flow on the interconnector exceeds 600MW about 2.3% of the time, and exceeds the level at the time of the November 1, 2015 event about 56% of the time.

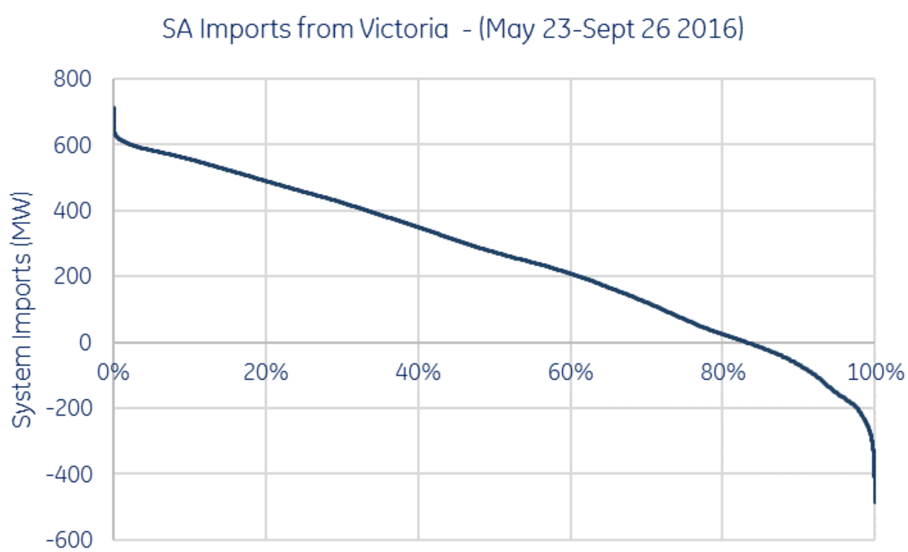

Figure 50 Recent (2016) Imports on Heywood Interconnector

6.2.1.1 Correlation between Inertia and Imports

Since the simultaneous condition of high import and low inertia represents the more challenging case, it is of interest to check on the correlation between high imports and low inertia. Figure 51 shows a scatter plot and correlation trend line for the data of the preceding two figures.

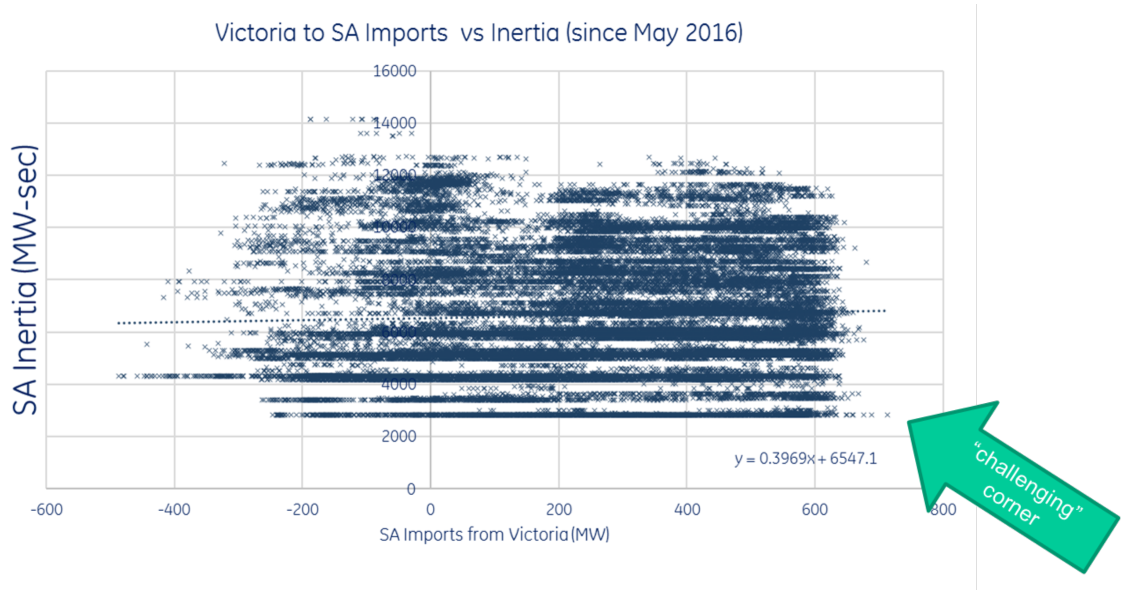

Figure 51 Inertia and Imports Correlation (2016)

The correlation is almost nil, with a slightly positive sign. The result indicates that no reasonable inference of system inertia can be made from the import flow levels. If mitigation is designed that is keyed to the inertia levels, other means to estimate system inertia will be required.

6.2.2 Context using example of November 1, 2015 event

For the purpose of understanding the requirements and benefits of providing FFR, it is useful to set aside the present autonomous underfrequency load-shedding scheme, and examine how the system would behave without it.

In Figure 52, the model is run for exactly the same initial conditions and event as the calibrated reference case, but with all the load-shedding – both voluntary and involuntary – blocked from acting. As would be expected, the system frequency swings much further, with the frequency nadir being lower and later.

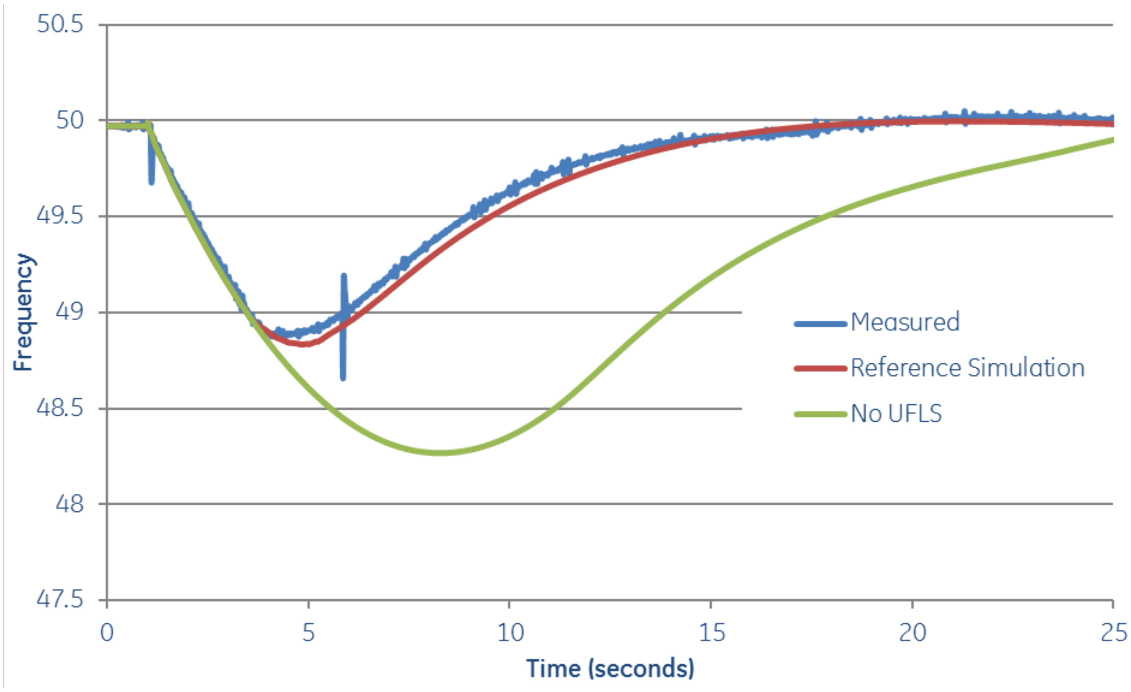

Figure 52 Simulation of November 1 Event without UFLS

Since a considerable amount of attention is given to inertia, it is instructive to consider, in isolation, the effect of altering inertia. Figure 53 presents a thought exercise, in which all other parameters are kept fixed, but the inertia is adjusted by modifying the MVA of synchronized unresponsive generation. Again, this is for the November 1, 2015 size event; i.e. tripping of the interconnector when loaded to 232MW.

The results are directionally as one would expect: for lower inertia, the initial RoCoF is higher, the frequency nadir occurs sooner and is deeper. And, conversely, adding inertia improves and delays the nadir. It is of some interest to note that even at greatly increased inertia, this particular event still results in UFLS. The orange curve is for inertia roughly twice the highest level presently observed in South Australia (per Figure 49 – maximum inertia with the exception of a few hours of outliers is around 12,000 MW-s).

The learning from this is useful, but limited: Committing generation for the sole purpose of *adding inertia* to improve the frequency nadir and avoid UFLS is not sufficient. If added commitment of synchronous generation adds to the available primary frequency response, then the results look different, as will be shown below.

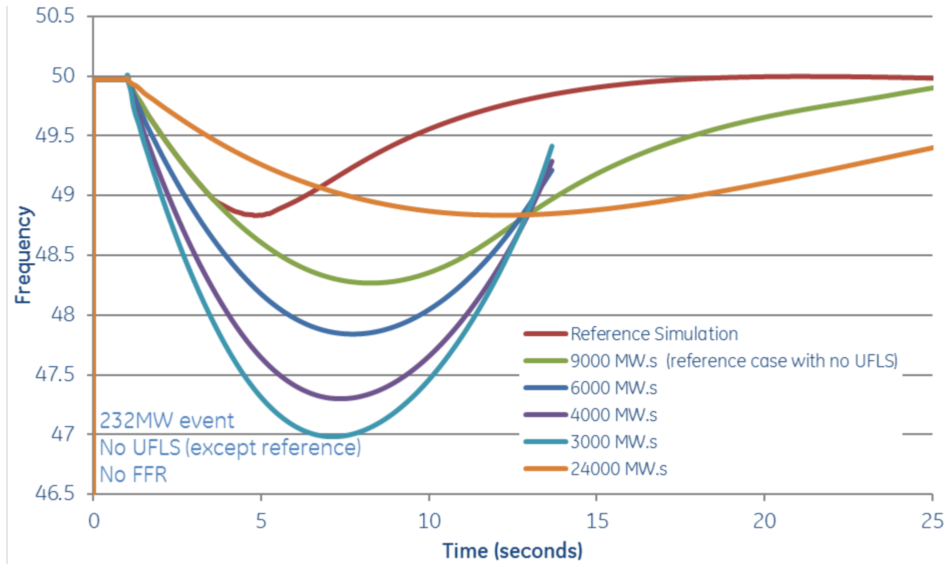

Figure 53 Frequency Response for Varying Inertia

These relationships are illuminating in two ways. They show:

1. The effect of changing inertia because economics drives different levels of inertia from synchronous generation in South Australia
2. The efficacy of deliberately changing (or limiting) inertia, outside of the economic dispatch and commitment. In other words, how effective is adding inertia compared to other means?

This shows that adding inertia alone is relatively ineffective means of avoiding UFLS. If the objective of adding inertia is to buy time for PFR and increase the likelihood that UFLS will successfully act to disconnect load to avoid dropping below 47Hz, then the reduction in RoCoF accompanying increased inertia will be beneficial.

The results of this sequence of cases can be mapped to show the relationships between system inertia and other performance metrics. Figure 54 shows how the nadir gets deeper with decreasing inertia (blue trace). If the amount of PFR is increased, all other things being equal, the nadir improves, as shown in the purple trace. In that sequence, the amount of PFR is doubled, by making twice as much South Australian thermal generation responsive. The time available to invoke other actions, besides UFLS, is shown with the green points: for this size event and operating condition, the threshold of UFLS at 49Hz is crossed in one second, whereas more than 2 seconds passes with inertia at 9000 MW-sec. The upper reaches of this curve are beyond the levels of inertia that South Australia sees today, but are illuminating in a historical context, as South Australia in recent years past would operate at these higher levels.

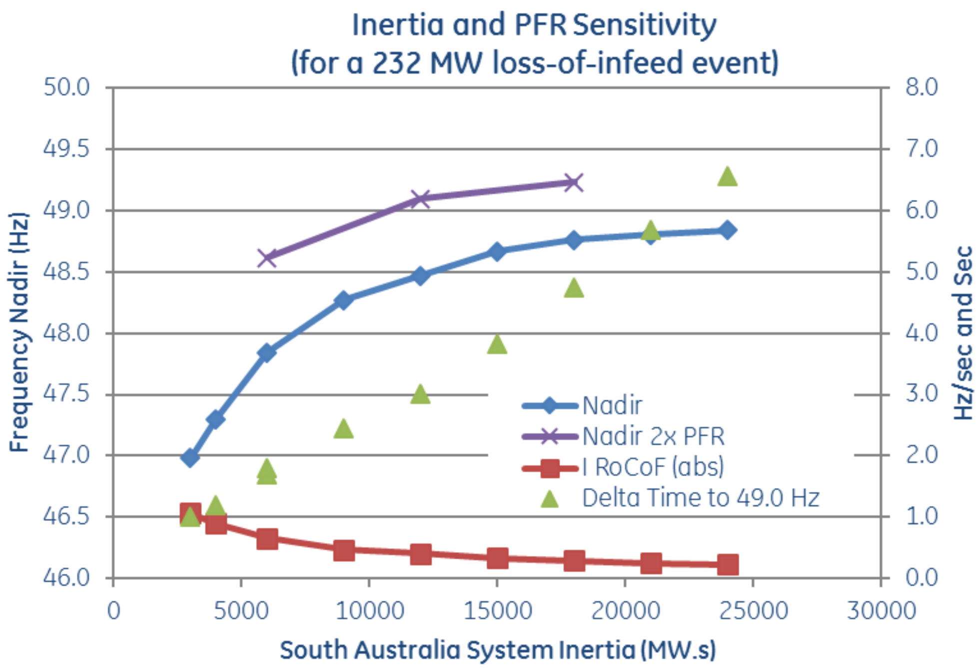

Figure 54 Inertia and Primary Frequency Response Relationships

This sequence of cases sheds some light on how the system behaves without FFR or UFLS. The primary focus of this work is to consider what can be done with FFR. For this discussion, we focus on avoiding any UFLS.

Figure 55 shows time traces of example cases in which discrete FFR action is taken at 0.5 seconds after the tie line trip event starts. The amount of PFR is held constant in this set of comparisons – in the sense that the units providing PFR are the same for each case and are at the same dispatch. FFR is generic here, in the form of power injected into the South Australian grid and sustained past the frequency nadir. The amount of FFR (i.e. the power level) is selected to make the frequency nadir 49Hz; that is the least amount necessary to avoid UFLS. As expected, the amount required decreases as inertia increases. In the upper trace, with inertia of 3000 MW-sec, the change in RoCoF is clearly visible when the FFR "hits" at about 49.4 Hz.

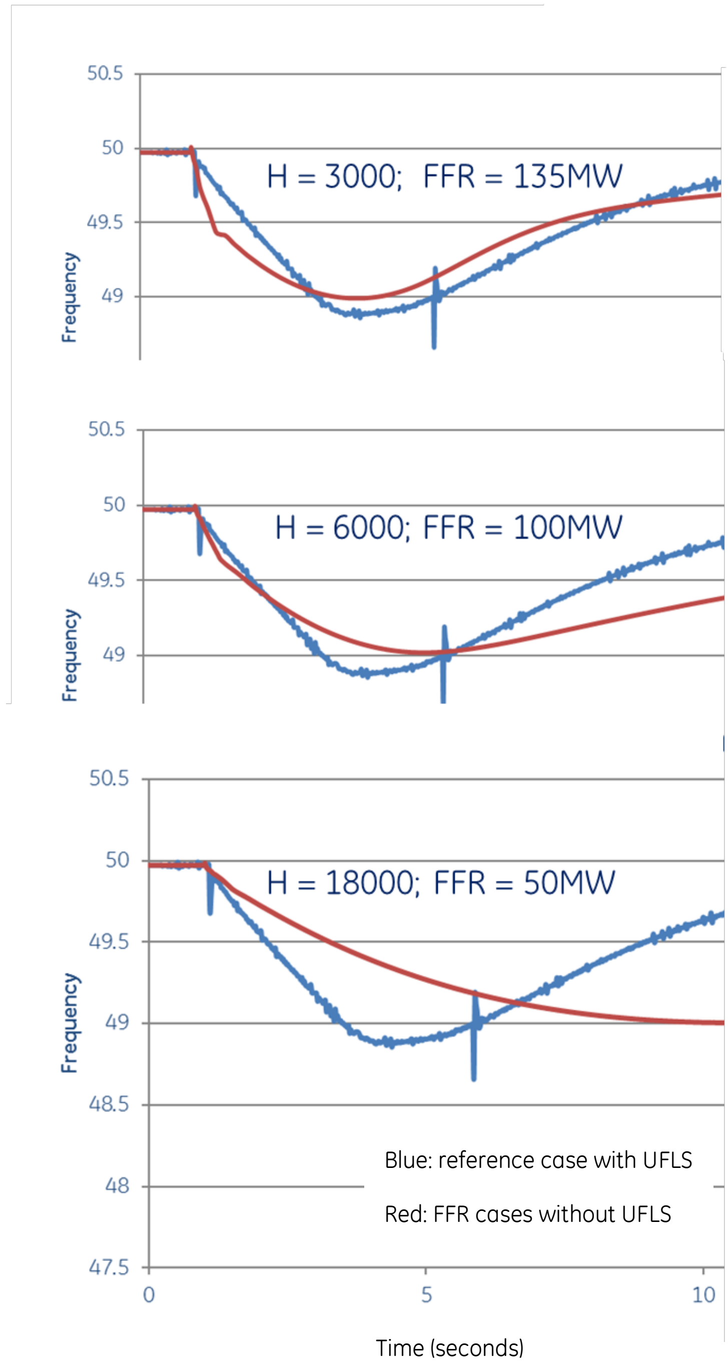

Figure 55 Illustration of FFR Requirement and Inertia

The results of this sequence of cases is presented in Figure 56. The red curve (which uses the right-hand axis) shows the power needed versus system inertia, and the blue curve shows the total energy injected between actuation and the frequency nadir. At first look, the fact that the required energy increases may be counter-intuitive, but the time until the frequency nadir is longer for the heavier system, and *more energy* is needed to arrest the decline of the heavier grid. Said differently, the lighter system requires a faster sharper application of arresting power, but it is easier to "catch" the lighter system.

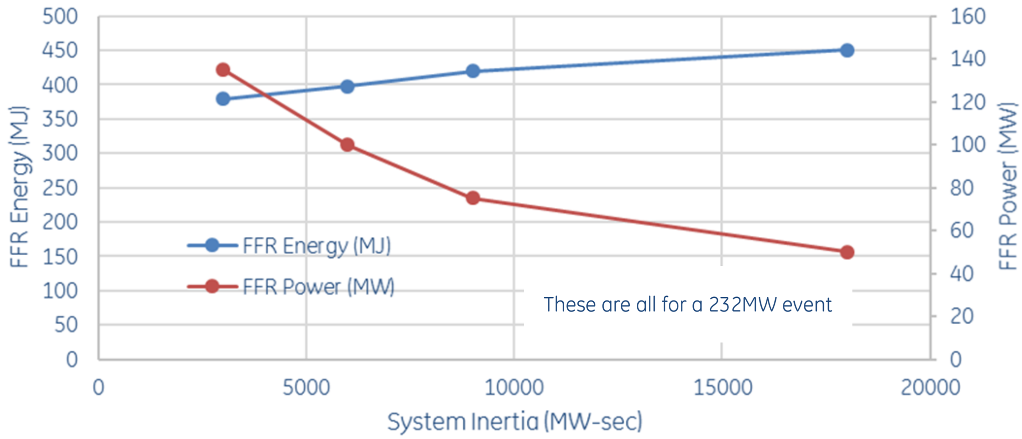

Figure 56 FFR Requirement to Avoid UFLS

This curve allows for comparisons of the efficacy of the FFR. First, we can consider the impact of FFR power rating on nadir, i.e. how much does the next incremental MW of "generic" FFR impact the frequency nadir? The blue trace in Figure 57 shows the marginal efficacy of the FFR. As discussed above, the FFR is more effective at arresting the decline of a lighter system. For example 1 MW of FFR produces 13mHz of improvement in the nadir (i.e. better margin to UFLS) at 3000MW-sec, compared to 9mHz of benefit at 9000MW-sec.

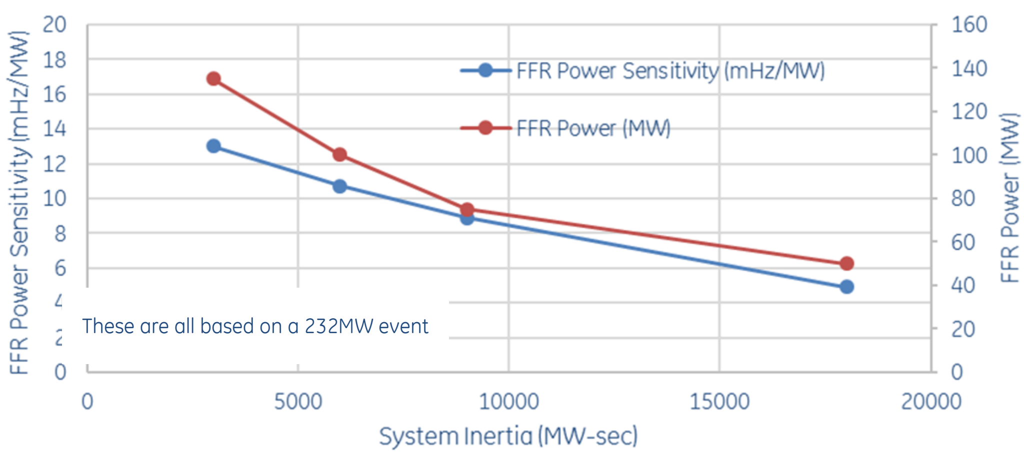

Figure 57 FFR Power Sensitivity

For a calibration exercise, in the figures it can be seen that, at inertia of 3000 MW-s (H = 3000), 130 MW of FFR is needed to avoid UFLS (for this 232 MW event). At H = 6000 MW-s, 75 MW is needed. This gives a basis for comparison: What does it cost to increase the inertia by 3000MW-s, compared to the cost of adding 55 MW of FFR? The answer to the 1st question is complex, and requires a production simulation to be answered with any fidelity. But, some calibration can be made. Consider 5 alternatives:

- 75MW of controlled voluntary load-shedding: cost is a few ones of A$M probably, but load shedding has societal costs.
- 75MW of IBFFR from wind turbines: inexpensive (but some cautions apply; see below)
- 75MW of BESS: say, A$75M, plus some variable operating costs.
- 3000MW-sec of flywheel or condenser: say, A$100M, plus some variable operating costs.
- 3000 MW-sec of forced commitment of existing generation: say 500MW at A$30/MWhr out-of-merit/uplift cost = A$15K/hour. For (say) 3000 hours/year = A$45M/annum. [Obviously this type of calculation MUST be done with a security constrained production simulation. Considerations for production cost simulations are discussed below]

We will revisit this exercise below, with consideration of extremely low inertia and a more generalized construct.

6.2.3 Sensitivity to Event Size (i.e. Initial Heywood Loading)

The results presented so far are based on a single large event: loss of 232MW of infeed. The question arises "How linear is this information?"

The next sequence of tests are against the event size; e.g. the loading on Heywood. In Figure 58, higher loading on Heywood is tested. The pair of cases are for the same inertia, no UFLS and the same PFR. The event is 50% bigger (~348MW), and RoCoF is roughly 50% worse, but the nadir is about twice as deep. Clearly, for this set of boundary conditions, the frequency excursion doesn't increase linearly with event size.

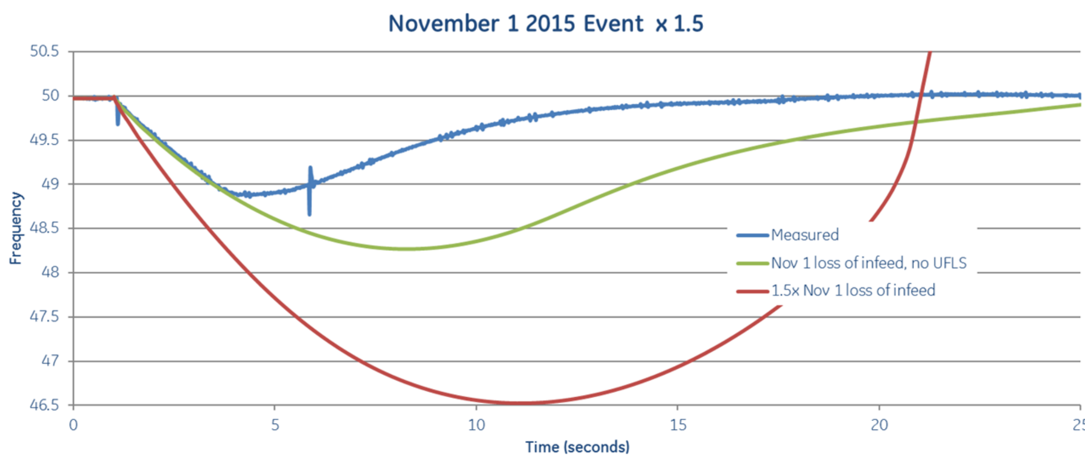

Figure 58 50% increase in event size

Closer inspection of the behavior of the system shows that the behavior *is* relatively linear, if viewed from the right perspective. In Figure 59, loci of frequency nadir versus the amount of FFR provided are shown. In these cases, the FFR is again a generic block switched at 0.5 seconds. The amount of FFR required to achieve a specific frequency nadir is very linear: the slope of the three traces is essentially the same. In Figure 60, for the same conditions, the minimum FFR required to avoid UFLS (i.e. keep the nadir up to 49Hz), is plotted. The amount of required FFR is very linear with the amount of loss-of-infeed. But extrapolation of the curve (dotted red) does NOT go through zero. This particular PFR "covers" approximately the first ~165MW of loss.

This test is for a fixed PFR and fixed H (9000 MW-sec) (and still no UFLS). The result suggests that for a given inertia (H) and PFR, the amount of required FFR is equal to infeed minus an amount of PFR "coverage". It is also of some interest to note that the time to frequency nadir shortens as the event gets bigger, but at these (now) relatively high inertias, the difference in timing is small. These results are quite linear, but simulations on a detailed, high fidelity model might show a higher degree of non-linearity.

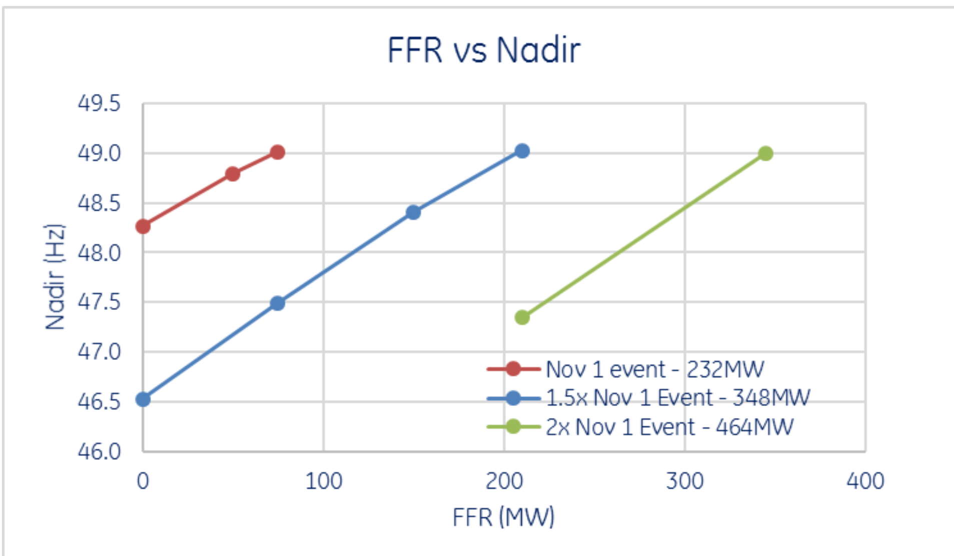

Figure 59 FFR vs. Frequency Nadir for different event size (fixed PFR and H)

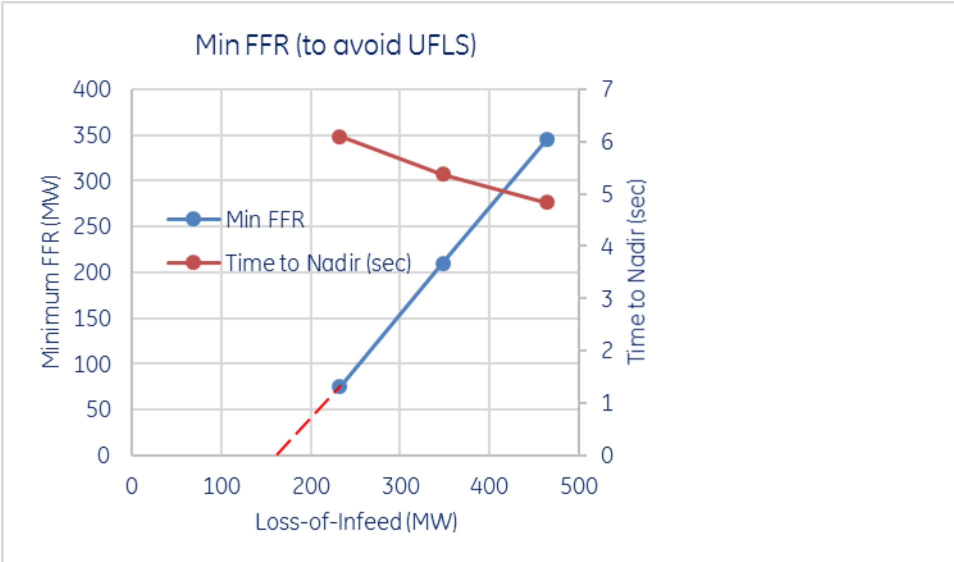

Figure 60 Minimum FFR necessary to Avoid UFLS (for one inertia and PFR condition)

6.2.4 Other Factors: Short Circuit Strength and Dynamic Performance of Synchronous Generation

The results presented above were deliberately constructed with minimal changes to other system characteristics such as inertia and PFR. But as the amount of synchronous generation drops, the contribution to system short circuit strength will drop as well. At the lower end of the range of inertia for this sequence of tests, the dynamic behavior of the voltage was observed to be rather poor. The model was tuned up for the condition, adding more power from wind and reducing the short circuit contribution from the synchronous machine equivalents. The equivalent exciter, and the voltage regulator on the equivalent wind plant were tuned as well. In Figure 61, a comparison of the same event (348MW loss) and inertia, with different dynamic models is shown. The newer case (red trace) has a lower frequency nadir (than the preceding case: green in this figure, red in Figure 58. This is mainly because the voltage was kept healthier during the swing. This may seem counter-intuitive, but the sensitivity of the load to voltage has been observed in our other work to occasionally be a significant factor for frequency stability investigations. Healthier voltage means more load active power, which exacerbates the generation-load imbalance. Nevertheless, healthier voltages are to be desired, and are an important part of maintaining stability.

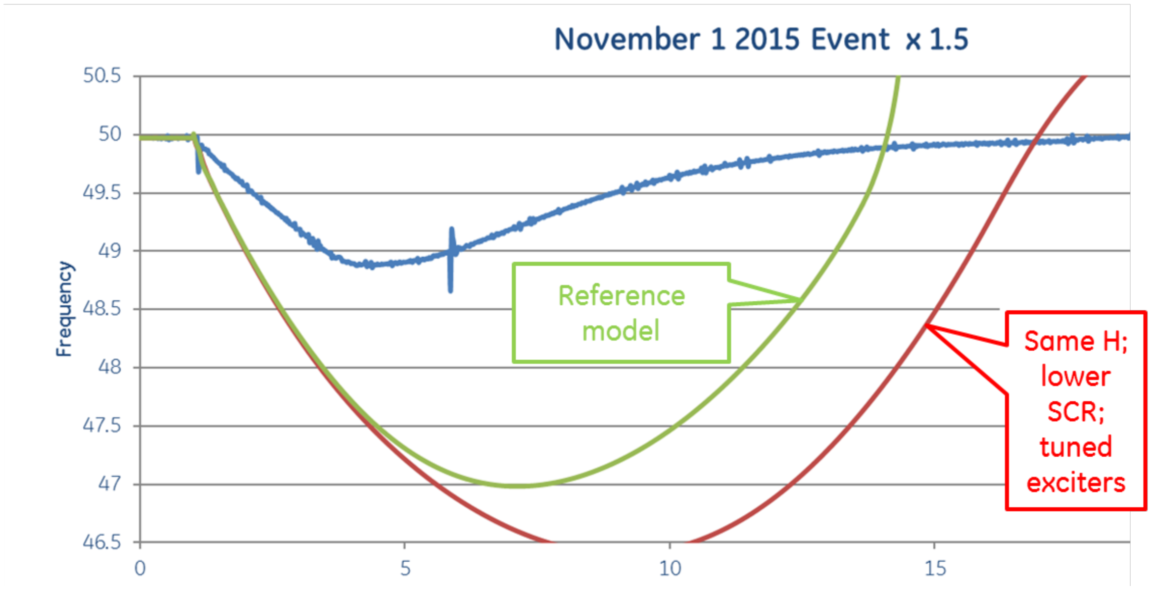

Figure 61 Effect of Other Parameters

It is nevertheless an important point, that the types of parametric illustrations provided here cannot be meaningfully made too precise. Modeling matters, as do details of dynamic behavior above and beyond frequency response.

The test to determine the FFR requirement to meet the 49Hz target is shown in Figure 62. It takes an extra 10 MW of FFR on top of the 135 MW required to avoid UFLS in the 348MW loss reference model case. The suggests, not surprisingly, that all calculations need some margin.

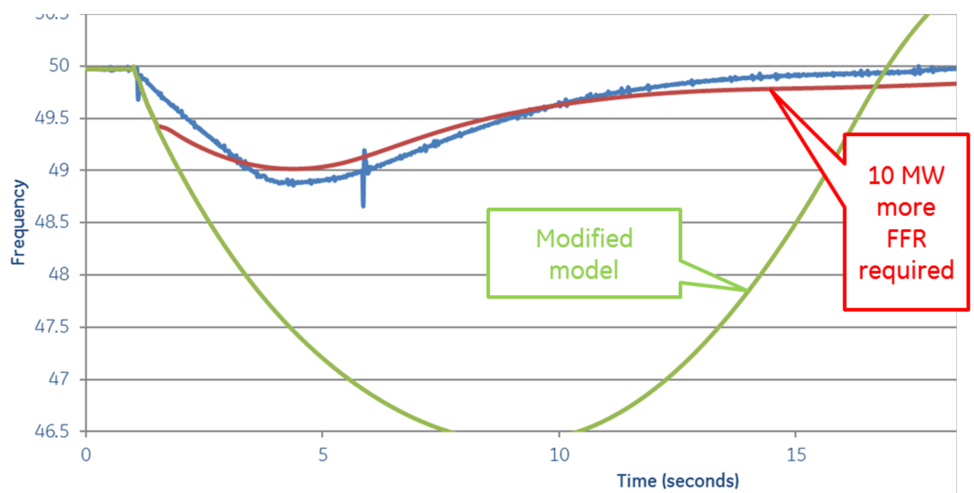

Figure 62 FFR Requirement with Modified Dynamic Model

6.2.5 How fast is fast enough?

In the discussions of section 2.2, we gave considerable attention to the challenge of rapid detection and decision making. But, clearly there is a need for speed when attempting to arrest the frequency decline for conditions of high power-load unbalance and low inertia.

Figure 63 shows the results of a sequence of cases with low system inertia (at 3000 MW-sec), subject to a 232 MW loss of infeed. The three cases are for different delay times between the initiating separation event and step actuation of 145MW of FFR. The case (red trace) for only 50ms represents that fastest time that we believe could conceivably be possible with the best available technology. On the other end, the green trace is for 500ms of delay, which represents a speed that is clearly possible with a range of different technologies and subject to a good level of confidence that the decision to trigger is proper. The blue trace is for ¼ second of delay, and represents an ambitious, but practically achievable speed of actuation. As might be expected, the very fast response gives the best performance, resulting in the highest frequency nadir. The ½ second response, with 145MW of FFR, is sufficient to meet the 49Hz target.

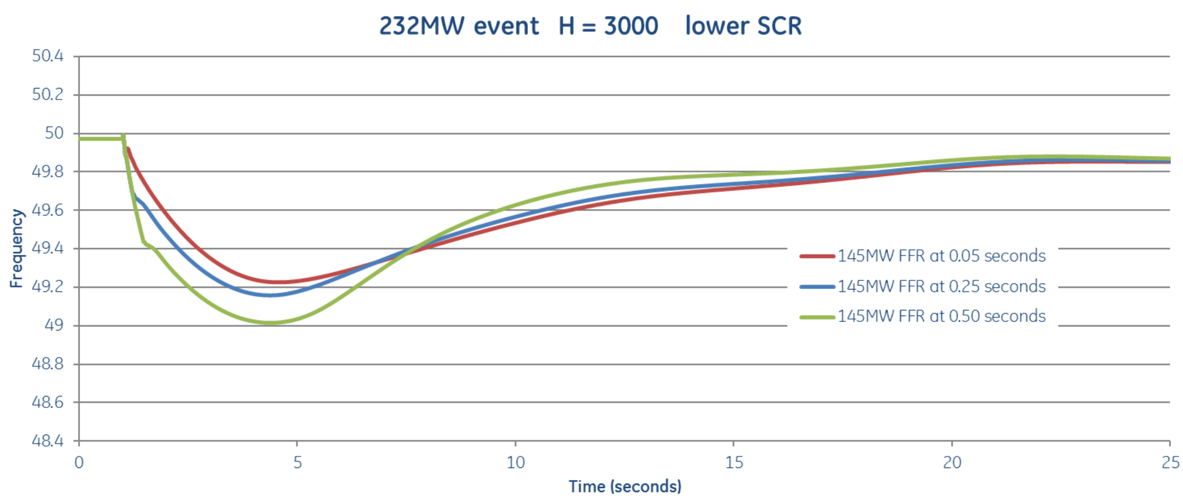

Figure 63 Efficacy vs Speed of FFR for Low Inertia

Figure 63 shows a diminishing return on speed, even for a "light" system. A faster FFR resource will require a lower power rating, primarily because it has more time to deliver arresting energy. Figure 64 shows a case (in red) in which the extremely (perhaps impractically) fast FFR power level is reduced to the minimum level required to avoid UFLS. In this case the required FFR power drops to 130MW from 145MW. So, triggering super-fast (50ms), saves 10% FFR power, over slow (500ms) triggering. As discussed in Section 2.2, there is higher risk of making an unnecessary trigger decision with shortened decision time.

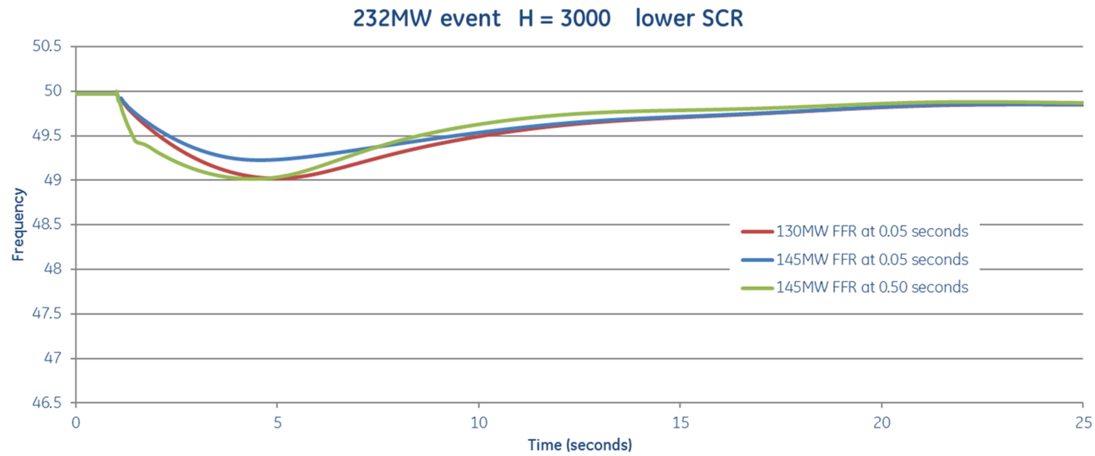

Figure 64 Speed vs Power for Minimum FFR (to avoid UFLS)

6.3 Requirements for South Australia: Future extremely low inertia system

The current trends in South Australia are towards even lower levels of inertia. In this section, the investigation is continued down to extremely low levels of inertia: 1000 MW-sec (which have been reported since this work began). This investigation deliberately does not extend below that level, even though it is fully conceivable that the present market and operational constructs could drive the system below 1000 MW-sec. Nevertheless, as discussed in Section 4.5, operation of large integrated power systems without any synchronous generation is presently the subject of a variety of research projects. Neither the available modeling tools (i.e. standard stability programs) nor presently offered inverter-based generation is suitable for this challenge today. Whether presently available technology, and indeed installed infrastructure, can be adapted to zero synchronous operation is an outstanding question.

In the tests of these very low inertia conditions, the MVA of wind generation is increased to allow most of the load not served by imports to be served by wind. The rating of the remaining generation continues to drop, so the drop in SCR and voltage support significantly affects the performance.

6.3.1 Stability Risks

A few simulations were performed that qualitatively highlight some risks as the system inertia drops to extremely low levels. In Figure 65, the exercise started in Figure 64 is continued. The case with inertia of 3000 MW-sec and 130 MW of FFR triggered at 50ms is shown in blue. As noted above, it is just sufficient to avoid UFLS. When the inertia drops to 2000 MW-sec (green trace), there is a slight decline in the frequency nadir. But when the inertia drops to 1000 MW-sec, the system suffers from a relatively dramatic stability failure. The failure is not simply due to an inadequate amount of FFR, as the case in Figure 66, with FFR increased to 200MW also fails.

On closer inspection, the very low H case fails were due, in significant part to poor voltage control and angular stability. The case was statically unstable, that is, even without a perturbation, the system could not operate as an island.

While the details of this rather simple simulation are not very quantitatively meaningful, this failure is almost certainly reflective of risk that will be observed in the real system: The dynamic behavior of the system with very low inertia and very low short circuit strength *will be* substantially different. Diligent checks of stability are absolutely required. A requirement for control improvements is likely.

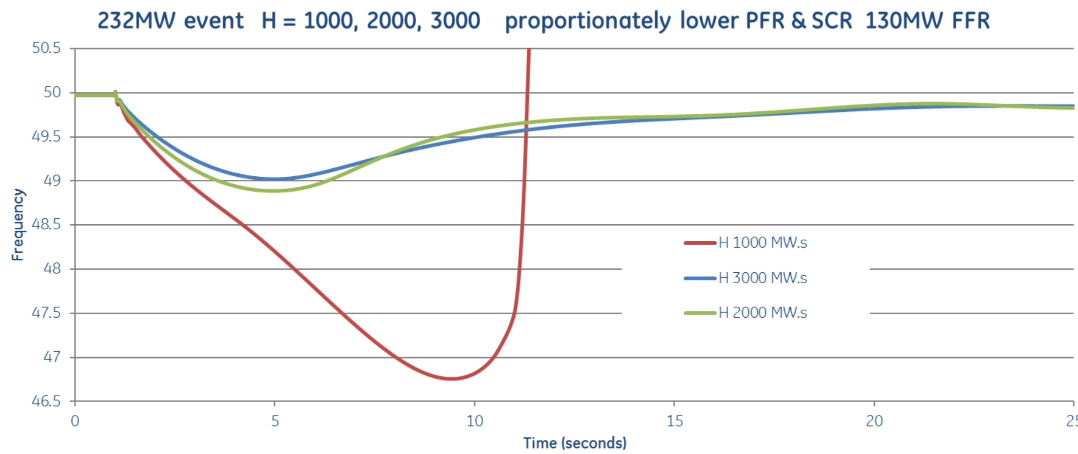

Figure 65 Stability Problem at Extremely Low Inertia - First Example

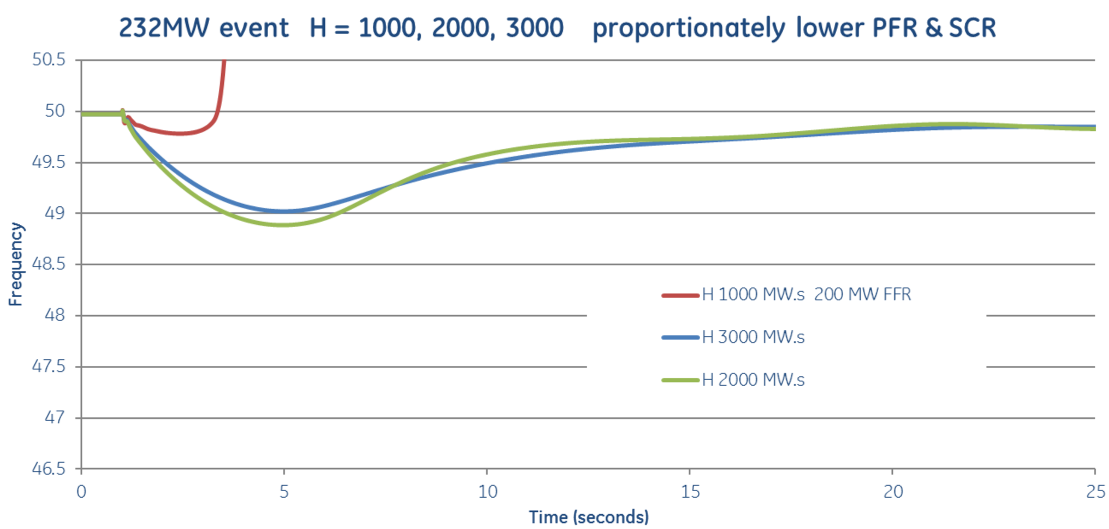

Figure 66 Stability Problem at Extremely Low Inertia - Second Example

6.3.2 Dynamics of extremely low inertia system

On closer inspection, the extremely low inertia case failures were due primarily to poor voltage control and angular stability. With improved tuning (but not extra equipment), the illustrative case can be made stable, for loss of the Heywood interconnector. Figure 67 shows a few successful trips of the Heywood interconnector for an inertia of 1000 MW-sec. In one case (the blue trace) the interconnector was unloaded. Trip of the line is therefore not a cause of power-load unbalance, but it is essential to understand that the operating condition *after* the line opens is hugely different than before. The effective gain on every regulator (both frequency and voltage control) is very different. There is no assurance that the system is stable in this condition, as indeed was the case for the system tested in the previous section. Opening the line causes a small, but not trivial perturbation in the system. In the other two cases, the loading of the tie line was 180 MW, and FFR equal to the tie line loading is applied. In the green trace, the FFR is the super-fast 50ms triggering, whereas the red trace is for triggering at ½ second. It is of interest to note that triggering an amount of FFR equal to pre-disturbance loading on interconnector results in a nearly instantaneous arrest of the frequency decline. Triggering ½ way to the UFLS threshold, still allows for successful recovery. The latitude for mismatch of FFR with event size declines with inertia.

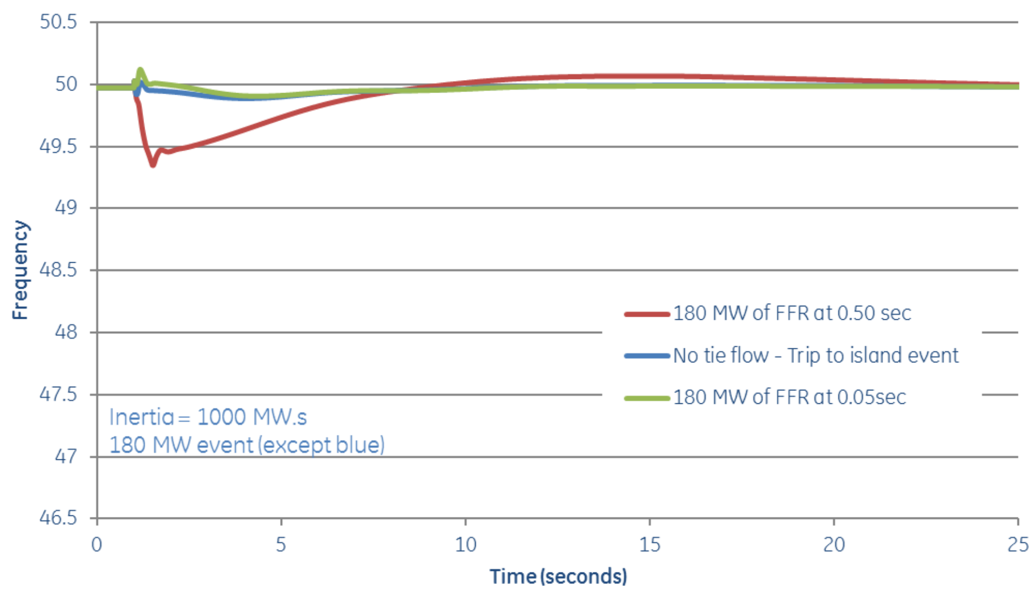

Figure 67 FFR equal to event size

6.3.3 Inertial-based FFR from Wind Generation

Before examining even larger events with extremely low inertia, we begin to look at the possible contribution of FFR with more sophisticated controls. That is, with closed-loop frequency sensitive controls, rather than the open-loop "triggered" behavior built into the cases presented to date. To that end, the cases presented here are based on the inertia-based FFR described at length in Section 3.2.

6.3.3.1 Tests of Wind IBFFR

An extension of the case with the lowest Inertia (H = 1000 MW-sec) is shown in Figure 68,. In this new case, the blue trace, the wind generation has an IBFRR feature enabled. In this case, there is 1000 MVA of wind generation producing about 720 MW of power. IBFFR is generally limited to about 10% of the wind turbine production, so was assumed that about 80MW of response will be available. Since this event is for 180MW of lost infeed, the case applies 100MW of switched FFR at 0.25 seconds.

The case has an interesting outcome: the frequency decline with IBFFR and 100 MW of FFR has a successful arrest of the initial frequency drop, with an initial frequency nadir of about 49.4Hz; i.e. about the same as for 180 MW of FFR. But, the withdrawal (and backlash) of the IBFFR is incompatible with PFR from the thermal plant. The primary frequency response does not "fill in" fast enough to allow the IBFFR to enter its recovery phase. The system frequency collapses. The case is a fail. For success, either the IBFFR needs to be sustained longer, or PFR from the thermal plant needs to be faster. Making IBFFR faster is not obviously very beneficial, which is consistent with earlier results. But, withdrawal could be delayed (within limits). We investigate further below.

This first test, "off the shelf" control settings were used with the IBFFR (labeled IBFFR#1 in the legend). In the case presented, control (red trace "IBFFR#2") is made considerably more aggressive and more sustained. The resultant frequency recovery is substantially better, but the case ultimately fails again during the recovery phase. PFR from the thermal plant still needs to be faster, and the result suggests that some tuning of PFR could coordinate with IBFFR better.

The point of these cases is to show that there are options and risks with closed-loop control. A further point is that IBFFR, and indeed many types of FFR, can have their behavior adjusted or tuned for the best performance. This observation applies for both closed-loop and open-loop controls. Broad questions of "does 'it' work?" are adequate to get guidance and calibration, but as with all the system dynamics under these high stress conditions, details of control matter greatly. This applies to design and application of any FFR, including that from wind plants.

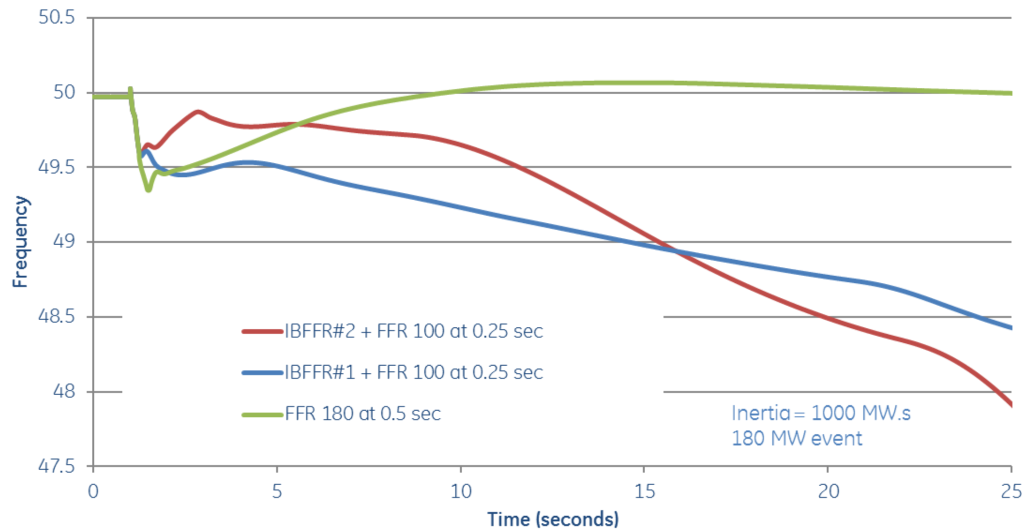

Figure 68 Comparison of 2 IBFFR Controls

6.3.3.2 Heywood at 450MW and Extremely Low Inertia

The cases presented in the previous section are for relatively modest levels of import. In those cases, the size of the event (180MW) was somewhat similar to the maximum contribution of the wind IBFRR – i.e. about 10% of production.

Here we examine an event that is large relative to available IBFFR from wind. The result of tripping the interconnector when loaded at 450MW are shown in Figure 69. There are three cases shown. In two of the cases, only switched FFR is applied. When exactly the amount of FFR (450MW) is triggered at 250ms (the red trace), the system survives. With a little less FFR (400MW), the frequency swing drops below 49Hz (green trace). Addition of IBFFR (the blue trace "WI") greatly helps the initial event, but poor coordination with PFR causes a fail. There are probably many ways to fix this, if there is adequate participation by the remaining PFR resources. Nevertheless, this highlights the narrowing window to have exactly the right amount of switched FFR triggered in a very short period. It is difficult, if not impossible, to imagine that autonomous, granular UFLS could possibly be relied upon to trip just the right amount of load. The addition of IBFFR clearly buys some benefit, and may allow other controls to take over. But, at least for this model, sustainability is a problem: the PFR needs to fill in faster.

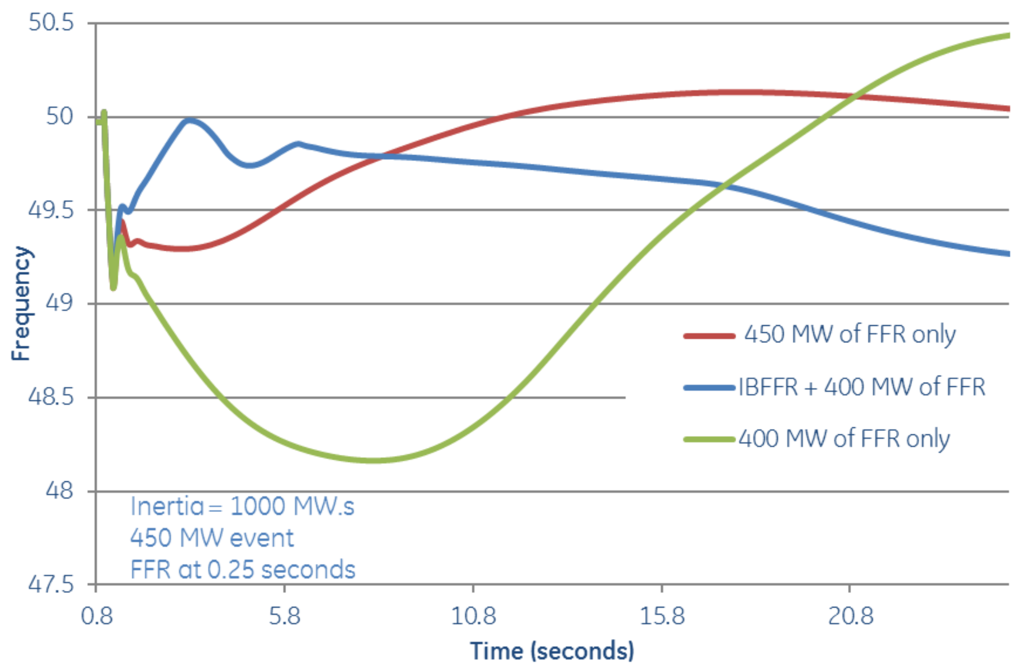

Figure 69 Illustration of Sensitivity to Balance at Low Inertia

In Figure 70, IBFFR is applied, combined with switched FFR equal to the lost infeed. In this case (the red trace), the IBFFR from wind helps improve and stabilize response with switched FFR. It does not "replace" the PFR (load balancing), but rather contributes significantly to the stability of the system during the initial swings. Open-loop IBFFR would not be as beneficial to stability. The contribution of the IBFFR could make otherwise insufficiently fast PFR "fast enough", although that is not the case in this specific example.

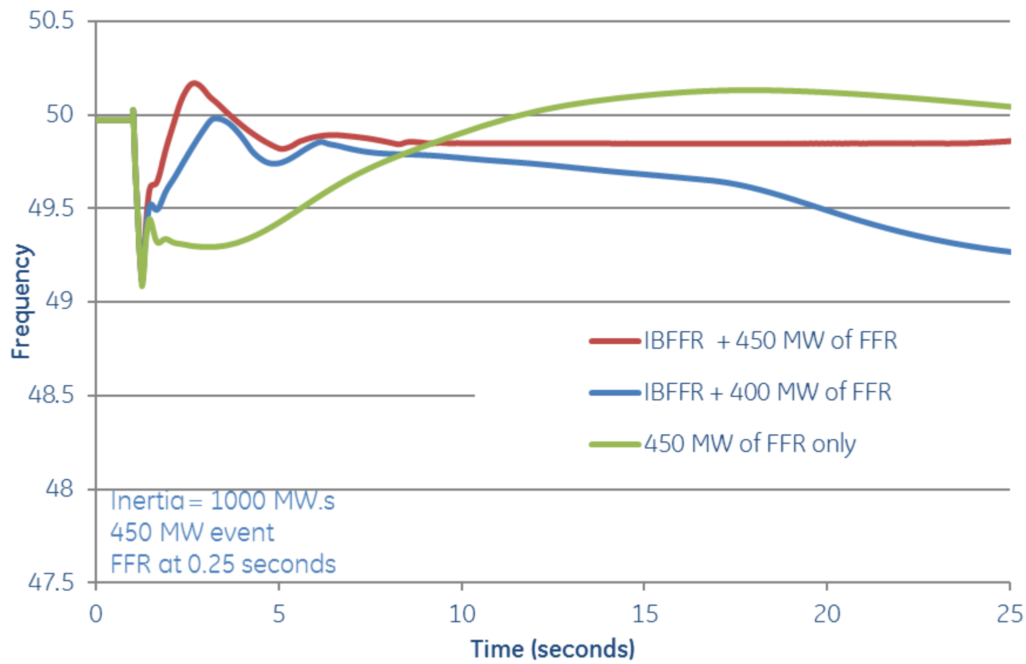

Figure 70 IBFFR benefits - rebalancing dynamics

6.3.3.3 Heywood at Limit and Extremely Low Inertia

A final case, in which we drive the system to the maximum stress under consideration in this report, is shown in Figure 71. In this case, the interconnector is loaded to 650 MW and system inertia is at 1000 MW-sec. The switched FFR alone (blue trace), triggered at ¼ second was very slightly late to avoid crossing the 49Hz threshold, but more important, the discrete action did not coordinate well with the PFR. The addition of IBFFR (red trace) helped stabilize this case, which by slightly speeding up the switched FFR to 200ms, avoids crossing the 49Hz threshold. The case suggests that, at best, relying solely on existing continuous controls and block FFR will not be robust. Note that this case, when run with a non-trivial fault before line trip was extremely unstable (not shown). Maintaining stability for this extremely stressed condition will be a non-trivial challenge. And of course, the extremely low SCR (out of scope here), will need to be considered.

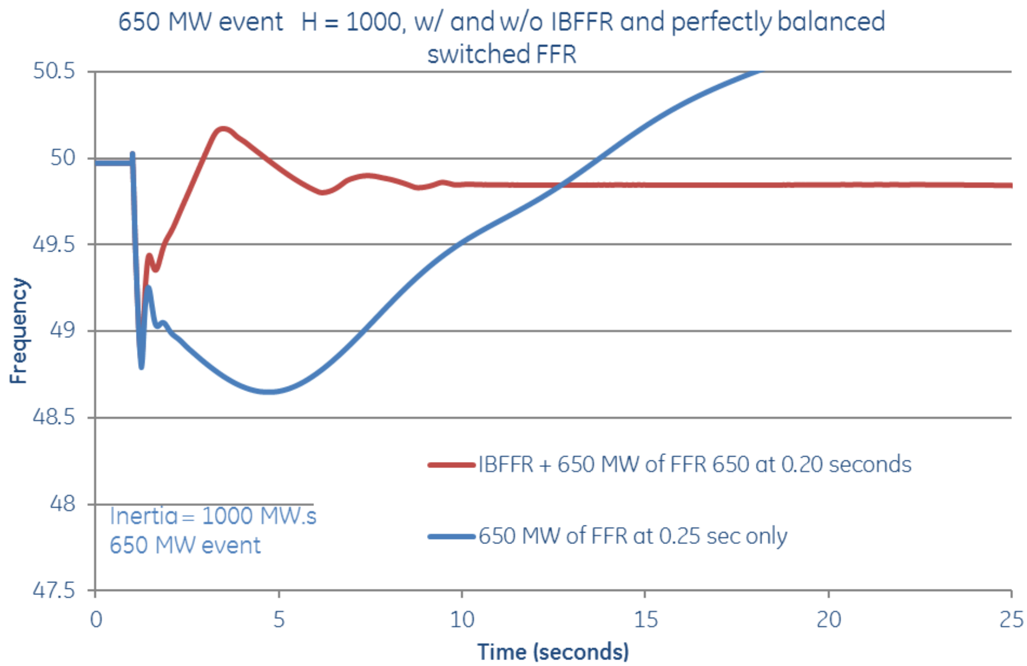

Figure 71 IBFFR at Limit of Low Inertia and High Heywood Loading

6.3.4 Summary: Trade-off of FFR requirement with Inertia

Results of this extremely low inertia investigation have been added to the results with moderate inertia (3000 MW-s and greater), and are shown in Figure 72. This figure ties together much of the essential findings of the lengthy discussions presented above.

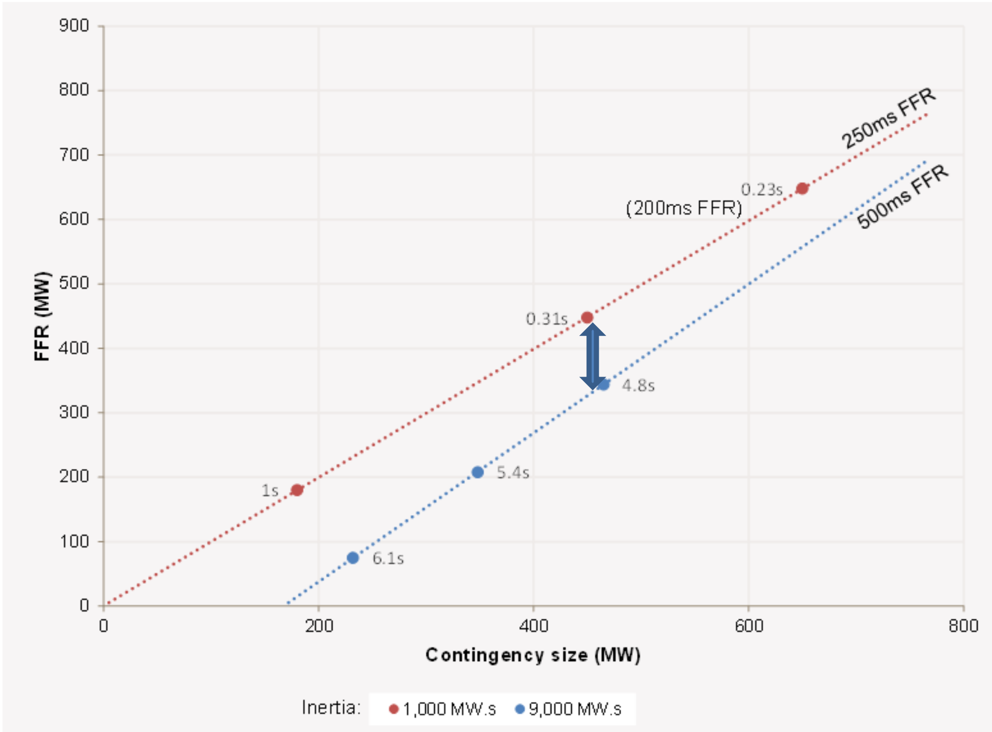

Figure 72 Relationship of Inertia to Min FFR and Event Size

In the figure, the efficacy of adding inertia relative to that of adding FFR is given by the difference between the red and blue curves. The time before the system reaches 49Hz without FFR is noted with each data point. Under these conditions, for example at 450MW import (as indicated by the blue arrow), it takes about 120MW more FFR to avoid UFLS when the system inertia is 1000 MW-sec than it does when the system inertia is 9000 MW-sec. Under these conditions, the FFR must act within about 250ms for the 1000 MW-sec condition, whereas this result for the 9000 MW-sec is based on 500ms action. The need for very rapid response disappears with higher inertia; this is the reason a 500ms action was used. The red trace, for very low inertia, is a very simple relationship: the FFR needs to equal the contingency size. This is because the PFR has so little time to act, that it essentially doesn't enter the relevant dynamics. The blue curve, with higher inertia, is displaced to the right. The slower RoCoF that goes with this higher inertia, buys time for the PFR to act. Effectively, the higher inertia creates the time necessary to get the equivalent of about 175MW of FFR from PFR. The slope of the curve is not the same as the red curve. This is because larger events, for the same inertia, result in faster frequency decline. The PFR is less effective, and therefore more FFR is required.

From a planning perspective, this allows some calibration: The capital cost to meet the marginal requirement for FFR is the difference in cost of (about) 330 MW of 0.5 second response FFR compared to the cost of 450 MW of 0.25 sec response FFR. A very rough estimate for 120MW dedicated energy storage device(s) would be A$120M. This figure is towards the high end of a range, since at least some of the necessary FFR could come from relatively less expensive resources, like IBFFR from wind generation. As noted above, the cost to *add* inertia (in the form of synchronous flywheels) is on the order of A$35/kw-sec, so 8000 MW-sec, would be roughly A$280M. Notice that this exercise assumes a degree of linearity that may be at odds with the exact details of the system performance. As has been noted (repeatedly), this simple model is insufficient to get precise results, nevertheless, this result (combined with the results shown in Figure 53) suggests that system solutions that focus solely on addition of inertia (in the form of flywheels or other additions that are solely evaluated on their inertia benefits) are unlikely to be economically optimal.

6.3.5 Determining Economic Benefit of FFR

To determine the relative marginal economic efficacy of adding FFR compared to imposing operating constraints in the form of constraints on South Australia inertia and imports from Victoria, detailed comparative production simulations are required. Unfortunately, imposing dynamic constraints on production simulations can be challenging. The results of dynamic simulations need to be mapped to static boundary conditions that can be modeled in a production simulation. There is industry precedent for this practice. For example, (in North America) transient stability derived path ratings are often mapped into (so-called) nomograms. Figure 72 a set of nomograms show how a production simulation could impose the interrelated constraint of import limit on the Heywood interconnector with committed synchronous inertia. The system is assumed to be constrained by “available response”. We introduce this concept here, because from a variable cost perspective (i.e. the constraints of a production simulation) willingness to shed load by UFLS (or other means) and the amount of FFR available are equivalent and additive. The different traces are measures of the sum of the allowed load shedding plus the amount of FFR provided. For clarity, consider the blue “no response” trace. If we suppose that no load shedding is allowed and that no resources are available to provide FFR, then this line represents the loading constraint on AC interface. At minimum inertia (of 1000 MW-s), zero power can be imported. As inertia increases (e.g. by commitment of synchronous generation in South Australia), then import limit eases, up to about 165MW at 9000MW-s. Now consider the red “200 MW” line. If 200 MW of FFR is available, then at minimum inertia, the AC import limit is 200MW. At 9000MW-s, import limit is 330MW. But, 200 MW of load shedding would produce the same result, as would (say) 100MW of FFR and 100MW of load shedding. This is a production costing constraint. Other measures, including those that might require capital expenditures (e.g. reactive compensation, or improved controls), could be needed, but they have no bearing on the variable cost of operation. Using pairs of simulations with this type of constraint allows a meaningful calculation or comparison of what FFR (or UFLS) is “worth”, from a variable cost of operations (and emissions) perspective. This approach would allow for the evaluation of added inertia as well; i.e. if dedicated flywheels were added to increase inertia all the time, the constraint would shift to the left, rather than up. A similar comparison of annual VOC would show the energy and emissions value.

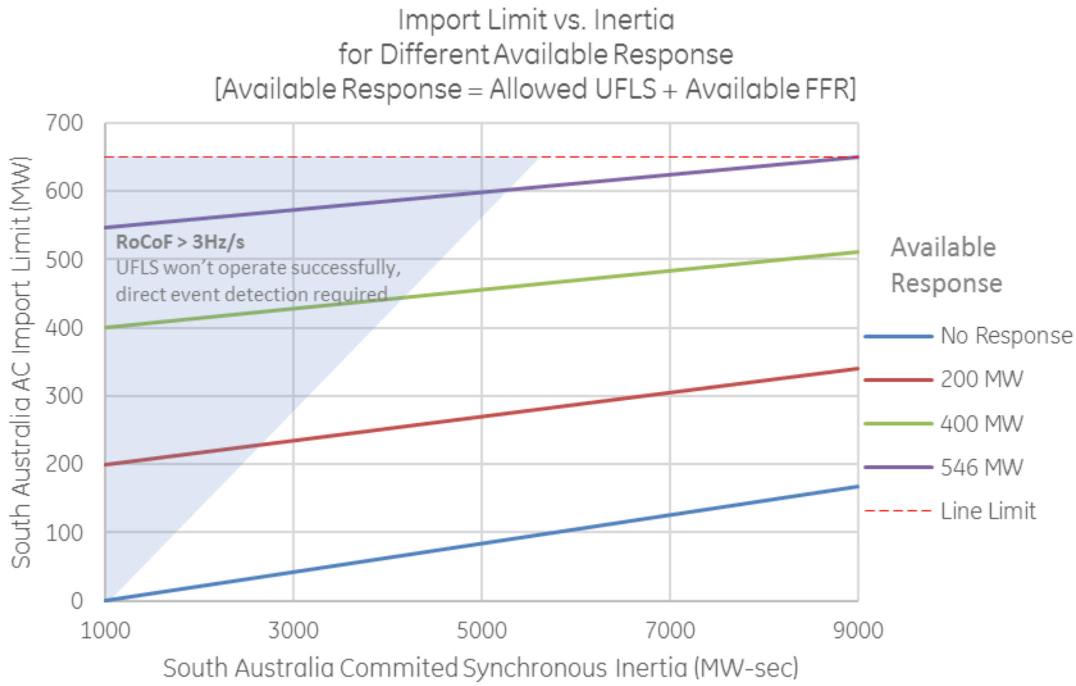

Figure 73 Production Cost Constraints for Import and South Australia Inertia

This exercise is, of necessity, isolating the effect of inertia and FFR. It is clear from the results presented above, that faster (more, better) primary frequency response would affect the amount of FFR needed for the 9000 MW-sec inertia case, whereas, faster conventional PFR response becomes less useful in avoiding UFLS for the 1000 MW-sec inertia case. Faster PFR, as for example that shown below (in Figure 74) from wind turbines will serve to offset load shedding. Also, there is a presumption that UFLS that acts at or below 49Hz will "save" the system. This means that the frequency swing should not go below 47Hz, and the system should not go black. For the low range of inertia, this almost certainly optimistic. That means that the region of the figure in the upper left hand side needs to be mostly, if not all, FFR rather than a mix of FFR and UFLS.

6.4 Wind with IBFFR and PFR control

The general context of this report is focused on options and requirements to provide FFR. However, as the results shown throughout this section reemphasize, the need and efficacy of FFR cannot be isolated from the amount and quality of primary frequency response available.

In Figure 74, a case in which the wind generation is dispatched (i.e. curtailed) to 90% of the power available from the wind, and primary frequency response is enabled. For clarity of understanding, it is important note that *for this illustration* we have assumed that wind speed is higher and have not actually reduced the output of the wind plant. The control settings for this case have less droop (higher frequency gain) than the typical 5% used (in North America) for synchronous generation. So, this is for the wind providing primary frequency response equal to 10% of the available wind. In this case, this corresponds to about 80MW of PFR available from the wind plants. The combination of the inertia-based FFR plus the primary response, allows for an equal (80 MW) reduction in the amount of block FFR (i.e. load tripping) necessary to stabilize the system. This illustrates two important points: (1) primary frequency response from wind plants is highly effective, and in this case, is substantially more effective than the PFR from the thermal plant, and (2) PFR from wind works well in combination with IBFFR. Again, as discussed in section 3.2.4.1, unlike the provision of IBFFR, provision of PFR from the wind plant has an opportunity cost for the generator equal to the value of the energy curtailed to enable the underfrequency response.

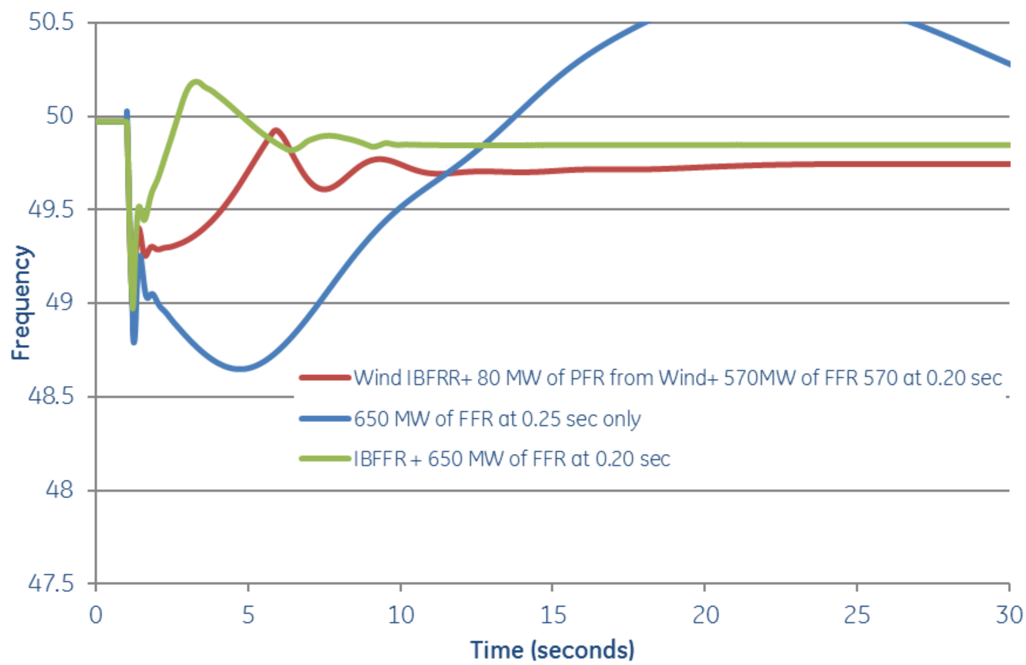

Figure 74 Illustration of Wind providing both IBFFR and Primary Frequency Response

6.4.1 Context to Other Systems

The size of the event under consideration in South Australia, i.e. up to 650 MW loss of infeed for complete loss of the Heywood Interconnector, is large relative to the system inertia in South Australia. Even at an inertia level of 9000 MW-s, the loss of 650 MW represents a ratio of about 15:1, inertia to event. At 1000 MW-s, the ratio is less than 2:1.

For some context, it is interesting to look at some other systems. For example, the design basis event for the US Western Interconnection (for which involuntary UFLS in not allowed) is about 3000MW. The minimum inertia is greater than 300,000 MW-s, so the ratio of inertia to event size is greater than 100:1.

In the isolated Texas system, the design basis event is 2750MW. In recent history, the lower end of the inertia range has a typically been around 120,000 MW-s, but ERCOT anticipates this to drop in the near future to around 70,000. In this case, the ratio of inertia:event size will be around 20:1.

The calculation is not as simple in Ireland, as there isn't a single event size. Rather they look at the worst loss of infeed that can happen for each event. Nevertheless, the 500MW DC infeeds tend to be the biggest limiting event. The system minimum inertia drops in the extreme to around 20,000MW-s. So, the minimum inertia:event ratio is on the order of 40:1.

The point of this is further reinforcement that this challenge is outside of the present industry experience for large systems.

6.5 Strategy for Extremes of Low Inertia and High Import

This review of technology and characteristics of the South Australia system suggests that a well-designed SPS could meet the necessary requirements of system security and economy. The severity of the problem to be addressed in South Australia combined with the substantial risk of overswing/backswing problems associated with discrete control actions almost certainly dictates that any discrete triggered FFR SPS continuously adapt to the exact level of power/load unbalance that will occur on separation of South Australia from the rest of the NEM.

The requirement for some sort of "master controller" that implements this function would periodically monitor the potential loss of infeed, i.e. the flows on the AC tie or across interfaces where the system can separate. The system would calculate the magnitude of the power/load imbalance in South Australia from separation events. This would be an indication of the magnitude of FFR and PFR response needed to stop the frequency decline and restore balance between generation and load. The master controller would also need to track the amount of FFR available from the various resources that could be triggered to act. The controller would then determine which FFR resources should be triggered, and possibly to what power level, if the ac intertie trips. This data would be updated every few seconds or minutes. Then, when the tie trips, the controller knows exactly what FFR devices are required to respond and the triggering signal can be sent out instantaneously. This is essentially a fast-acting control scheme where the triggered action is continuously updated in advance to adapt to changing grid operating conditions.

There is precedent for this type of situationally adjusted arming. For example, a similar adaptation is used to arm the remedial action scheme (RAS, aka SPS) on the Pacific HVDC Intertie in the western US. Also, the Intermountain HVDC system in the US was equipped with a similar "contingency arming scheme" scheme that could adapt control actions to ten different monitored system events.

The impetus for this type of “adaptive” triggered control action is because, as the illustrative analysis above showed, the magnitude of the control action needs to be closely matched to the magnitude of the disturbance. If the control action is too small, frequency would decline too far and load shedding could be triggered. On the other hand, if the control action is too large, the imbalance of generation and load would change sign and frequency would quickly rise and overshoot 50Hz. An underfrequency event would inadvertently cause an overfrequency event. By adapting the magnitude of triggered FFR to system operating conditions, the frequency decline can be quickly arrested in a manner that will allow PFR resources to handle the remaining unbalance between generation and load. The inclusion of at least some fast acting closed-loop frequency sensitive FFR (including IBFFR) will increase the robustness of the system, and aid maintaining stability.

6.5.1 Special Protection Scheme summary

Specifically, in summary, we recommend consideration of an SPS that:

- Is *primarily* based on fast load tripping,
- Detects islanding quickly,
- Is armed *a priori* to rapidly trip load “equal” to AC imports
- Is augmented by other fast technologies

Some key summary points are:

- For all but the very most extreme conditions, 250ms for action is fast enough to avoid involuntary UFLS.
- The consensus of (GE) experts is that a secure system with this speed is possible. Faster *might* be possible.
- The requirement for exactness of “equal” increases as RoCoF increases, which occurs as:
  - import increases
  - inertia in South Australia decreases
- Other fast technologies will help stabilize and improve robustness of such a scheme.

6.5.1.1 An SPS won’t be easy

There are significant challenges and functions of an SPS. To determine how much FFR, i.e. load tripping plus other FFR must be armed at any instant of operation is complicated by:

What is left after separation?

- Where does the system break?
- It isn’t just Heywood at the South Australia-Victoria border; there are many places that the South Australia system can break away from the rest of the NEM.

Distributed PV?

- Are you tripping generation, too?
- How much do you get by tripping?

6.5.1.2 The SPS may be needed for other risks

It is likely that the SPS can provide benefits for other classes of events. Most notably:

- Overload scenarios of the general type that occurred on Sept 28, 2016 may be addressed, at least in part, by such a scheme.
- The required amount of load tripping (and other FFR actuation) will not be the same as for complete loss of in-feed, but will likely have some similar dependencies.
- Such an SPS might also be useful for managing partial loss of the interconnection as well. Detailed design of the SPS should consider such other functions.

6.5.2 Role of other FFR technologies

The SPS must distinguish between Sustained and Transient FFR technologies:

- Sustained FFR provides power until secondary rebalancing (e.g. unit start-up, resync, manual load adjustments, etc.). It is a substitute for PFR.
- Transient FFR "buys time" for slower PFR to act. (AEMO's present 6 second product).
- Provision of PFR from resources other than synchronous generation should be considered. An example of wind generation providing PFR shows good performance.

Technologies discussed in this report that provide a sustained response (other than load shedding) include:

- BESS, curtailed solar & wind, possibly some other energy storage technologies and possibly the HVDC.
- Each MW of these technologies displaces 1 or more MW of load shed.

Technologies discussed in this report that provide a transient response include:

- Includes inertia-based FFR from Wind turbines, high power/lower energy storage (high inertia synchronous condensers, ultra-caps). A short duration HVDC FFR function is more reasonably regarded as a transient resource.

6.5.3 Requirements on other FFR Technologies

Speed of response

- Devices should be able to respond quickly.
- Analysis suggests that extremely fast response (e.g. less than 100ms) is not required, and may not produce incremental improvement over responses on the order ¼ second (even slower for the present system with 3000 MW-sec inertia).
- Closed loop response; with minimal deadband and over-frequency response, will be highly desirable. Control should be stable and well behaved for backswing over-frequencies.

Transient FFR

- Careful consideration to the timing of withdrawal and recovery (if any; e.g. for inertia-based FFR for wind turbines) is needed.
- For IBFFR, response should begin to be withdrawn at the frequency nadir; but not too quickly (assuming that response can be sustained that long).
- Management of component duties within the HVDC will be needed in the event that FFR is designed to temporarily overdrive the system. These are likely to be the limiting factor in FFR from this resource.

6.6 Autonomous Schemes and Backswing Risk

Autonomous, distributed UFLS is difficult to get exactly right in most systems. Finely grained and staggered UFLS are the standard approach to attempting to have the right amount of load trip for different events. Such a system, with many small steps, separated by a combination of different delays and frequency thresholds, requires time to self-coordinate.

With the potentially very wide range of inertia and RoCoF, it will be even more difficult to get right in South Australia…at least using local measurements. A system that arms autonomous, distributed UFLS based on current (real-time) risk (i.e. loading of Heywood), would help. This is a similar concept to that outlined above, but with the distinction that arming would be set by a centralized system, but triggering would be based on local measurements. This approach is, in some regards simpler and probably less expensive. However, since there is a multiplicity of separation scenarios, such arming is necessarily approximate, since local measurements, unlike a complete SPS, will be unable to differentiate different events.

This inevitable imperfection raises the risk of tripping too much load, causing a backswing of frequency to over 50Hz. Around the world, backswing has caused many system-wide blackouts. In Figure 75, an illustration of the risk is shown. Note that this just a drawing, and not a simulation. The call-outs point out the sequence. New AS4777 may help mitigate risk of backswing by reducing power generation by distributed resources, but it also has risk of exacerbating loss of generation. Note that this is not *necessarily* a problem, but rather that consideration of the contribution of AS4777 compliant resources should be included in simulations intended to evaluate the risk. The big unknown and risk in this scenario is generation tripping during the high frequency. Causes could include lean blown out (LBO) of gas turbines, action of protection, or sensitivity to high positive RoCoF. The point of this illustration is that too much open-loop response, including excessive UFLS, is a significant risk. And conversely, fast continuously acting resources will add robustness and reduce the requirement for precision in tripping just the right amount of load.

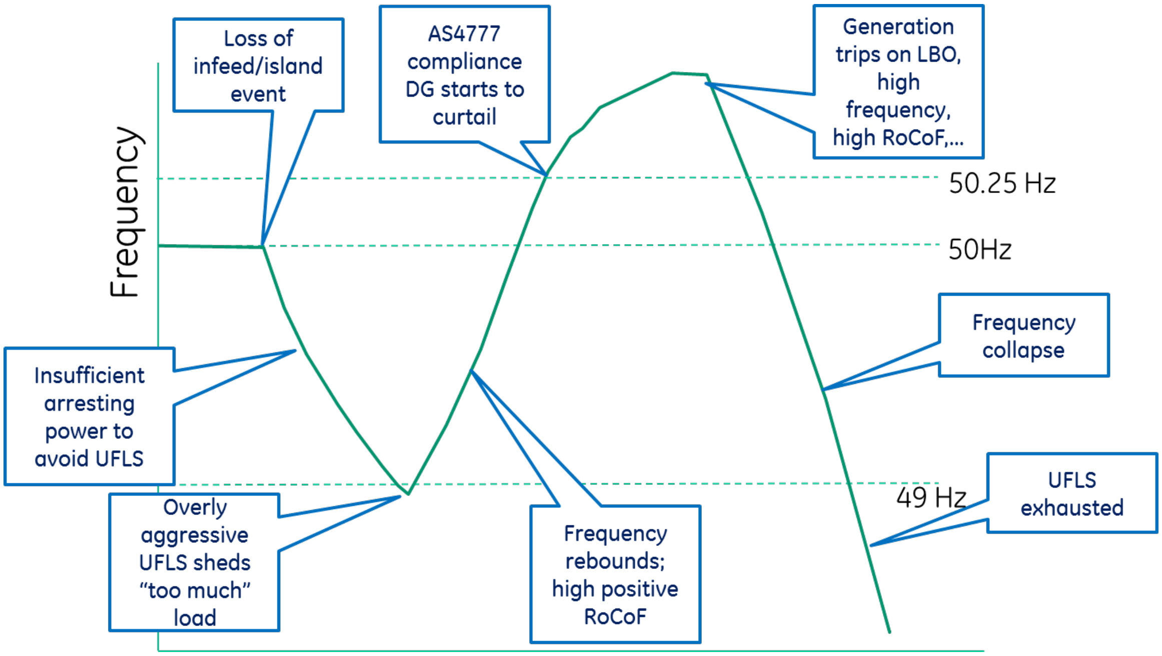

Figure 75 Backswing Risk Illustration

7 FFR AND SIR ANCILLARY SERVICE SPECIFICATION

In this section, we present an example of the mechanism by which AEMO can determine the necessary amount of the new FFR ancillary service to procure. The intent of the exercise is to show the relationship between synchronous inertia, the existing (or modified) primary frequency response (PFR) ancillary service and a new FFR ancillary service. As shown in the preceding section, the physics of the problem dictates that these three aspects are inexorably intertwined, and the most effective and economic procurement of ancillary services will need to consider all three.

The present nascent industry practice for FFR products has created a somewhat limited definition of FFR that tends to limit or exclude technologies that could be effective in meeting AEMO's objectives (and in opposition to one of the stated objectives of the statement of work: "...specification should facilitate the broadest possible participation...")

To facilitate procurement from the widest range of technologies, we illustrate a process by which FFR resources with a range of response characteristics can be normalized to a reference FFR functionality, (e.g. a block of arresting energy). This provides a mechanism by which different characteristic resources can bid and be compensated for equivalent beneficial impact on system frequency performance. This concept is an extension of the basic underlying SIR-FFR-PFR trichotomy.

The idea is to create relationships that can be used in the co-optimized market clearing of energy and ancillary services. The underlying assumption is that introduction of too many ancillary services will be hard to police and to achieve liquidity in the market. Thus, the expectation is that multiple characteristic resources will be able to bid into a single (or few, if necessary) products, and be "normalized" in their systemic value to AEMO.

7.1 Power, Energy and Timing

Most of the systems investigating FFR, have tended to have rather strict definitions of the expected behavior of FFR. This has the potential to limit the technology choices, since strict adherence to proscribed power signatures may exclude some technologies. Further, it is by no means assured that the prescribed performances provide the most benefit to the grid.

In this context, the questions of "what behaviors produce the best outcomes?" and "how does one compare resources with very different behavior characteristics in their abilities to maintain system security?" become pressing. The concepts and results presented in this section present a broader framework for characterizing the quantitative impact on system security, particularly frequency nadir, for any resource that contributes arresting power. The intent is to allow "apples-to-apples" comparison between resources, so that compensation, via markets or other means, can confidently:

- Reward the behavior that best meets the reliability need,
- Allow the broadest participation by resources that can contribute,
- Ensure that the right amount of resources is procured, and

- Pay fairly for the benefits delivered.

7.2 Evaluation of FFR Performance

In the following sequence, we present a construct for evaluation of FFR. To that end, the performance of the South Australia illustrative system without FFR is established as a baseline. The baseline case is for a relatively high level of system inertia (12,000 MW-s), with the intent of illustrating an approach that will be useful across all of the NEM and not limited to just South Australia.

A sequence of tests, in which a 250MW, 250ms "block" of idealized FFR energy is injected at successive ½ second intervals (the total energy injected is therefore 62.5 MJ). The concept is illustrated in Figure 76, with the baseline shown as the blue trace, and one perturbation. The perturbation is the red trace, in which the energy pulse starts ½ second (500ms) after the event starts and ends after another ¼ (250ms).

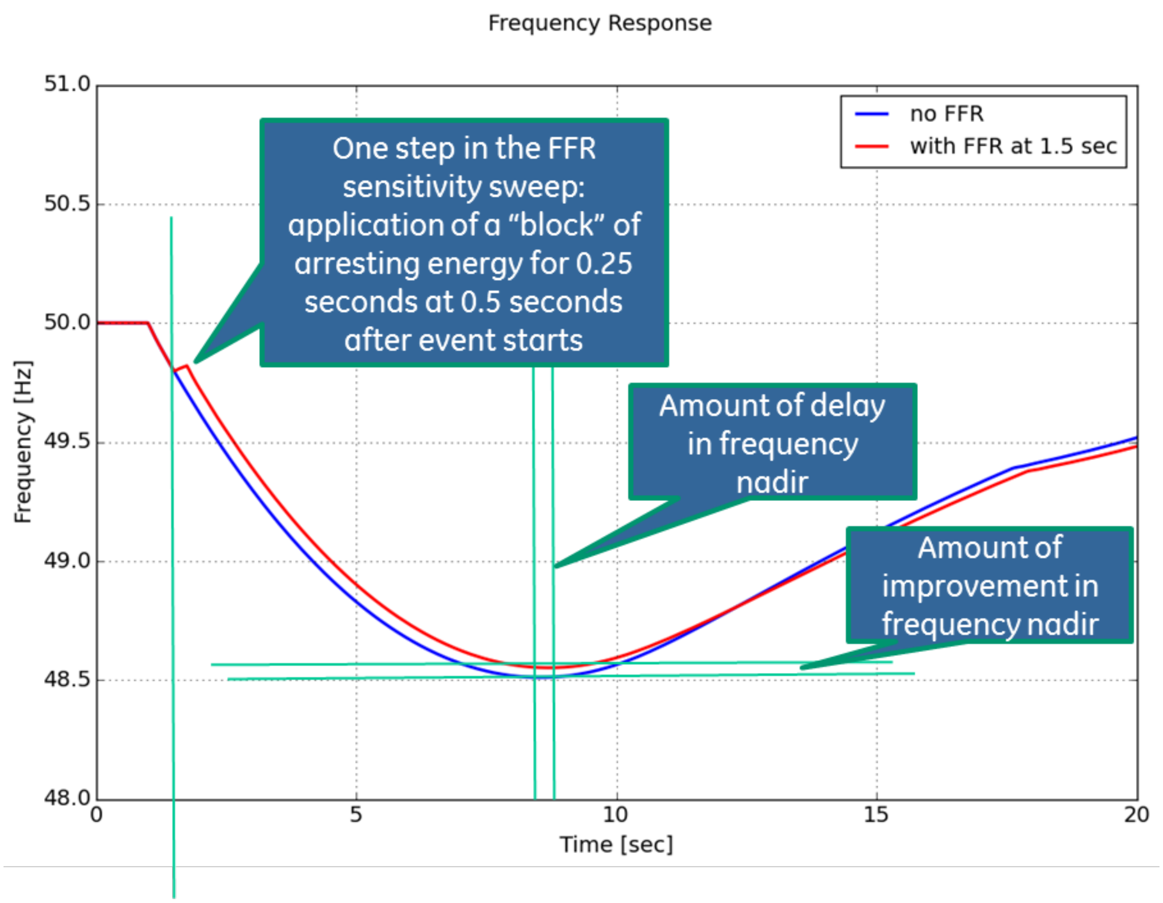

Figure 76 "Ideal" FFR perturbation test

The added FFR arresting energy improves the frequency nadir, and delays the time at which the nadir occurs. This occurs, to differing degrees, for each interval up to the time of the nadir. Obviously, FFR injections that occur after the nadir have no benefit (arguably by definition). A composite of all the cases, run in a sequence up to four seconds into the event is shown in Figure 77.

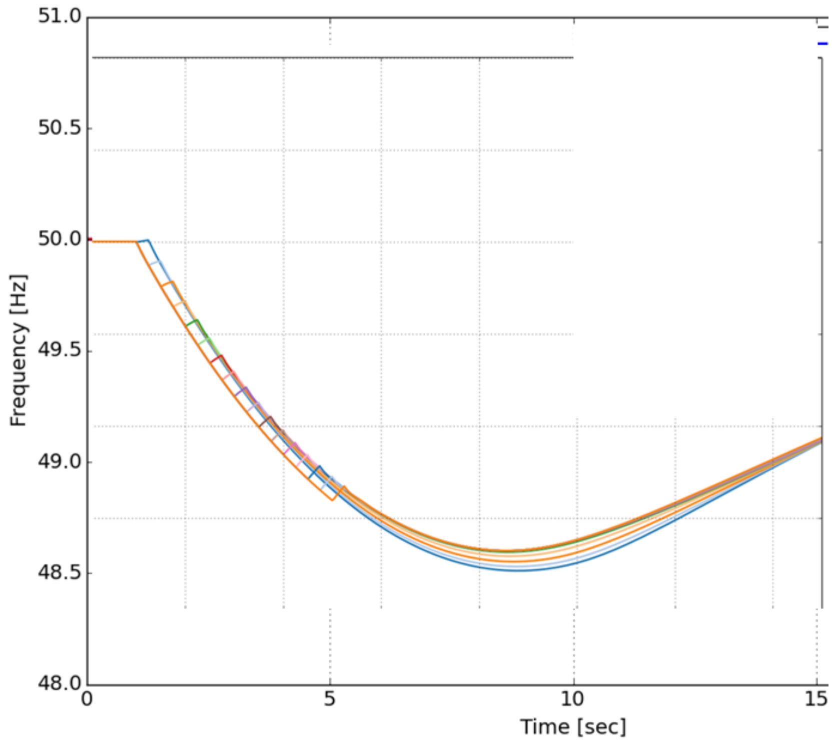

Figure 77 Frequency Response for Trip of Heywood Interconnector

7.2.1 FFR Response Efficacy Mapping

The sequence of cases results in differing impact on the frequency nadir, even though the arresting power and energy are exactly the same for each perturbation. The change in nadir is shown in Figure 78.

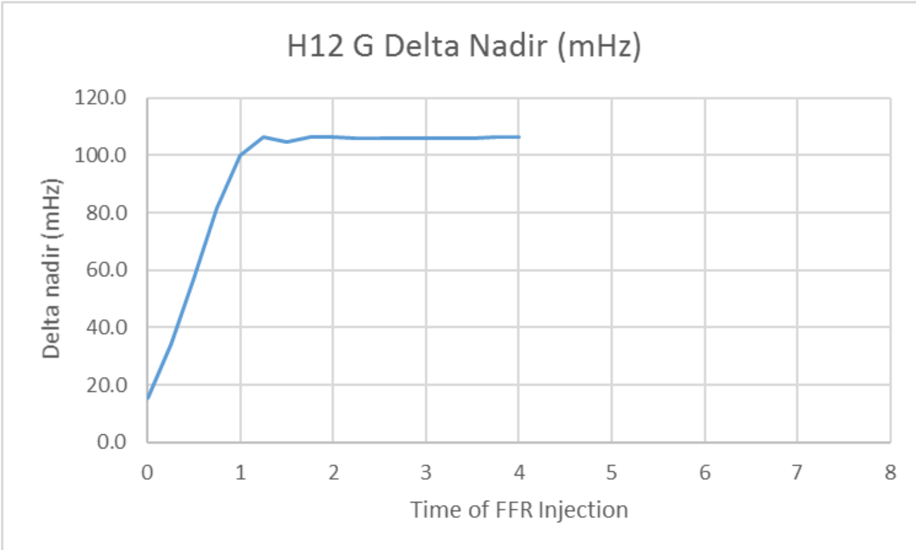

Figure 78 Nadir Improvement vs. Timing of Arresting Power Injection

This result is then linearly mapped, or normalized to reflect how any arresting energy applied at the given time will impact the frequency nadir, as shown in Figure 79. (The green is extrapolated data; the efficacy tends to roll off as the nadir approaches, partly because there is less scope for improvement as the frequency approaches the nadir. The extrapolation is unimportant for the point of this illustration. When the transfer function is actually calculated with a complete system model, the perturbation simulations should be run to times past the nadir.)

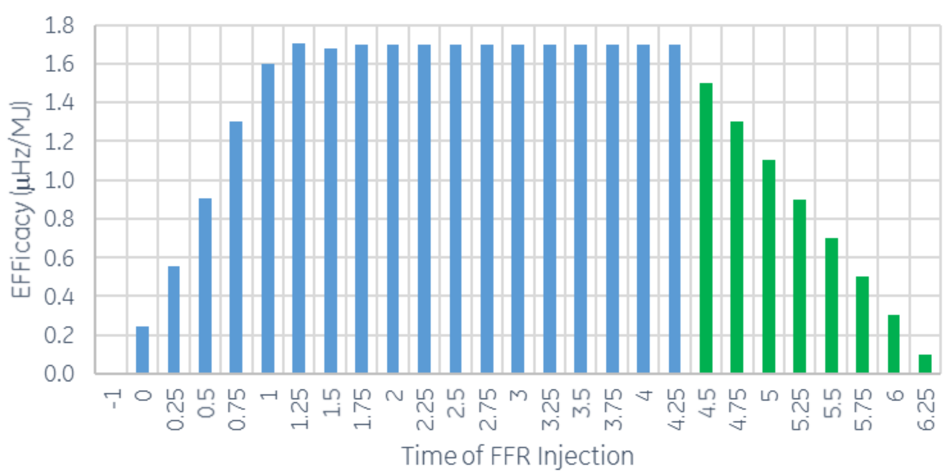

Figure 79 FFR Efficacy Mapping

It is significant to note that injection of arresting energy later in the event has much higher efficacy at improving the nadir compared to injection that occurs immediately after the event starts. This is due to the dynamic interaction of the FFR with the PFR. By allowing a more severe initial RoCoF, the PFR controls act more aggressively and produce better outcomes.

This sensitivity rather dramatically shows that faster response of FFR is not necessarily better. In this case, injections after 1 second are more effective. The authors have tested this approach on larger, more complex systems, and have found this trend. Broadly, we have seen that in systems with frequency nadirs occurring several seconds after the event have the best efficacy about one second before the frequency nadir. The behavior for a system with very fast nadirs will have different quantitative results.

This may all seem rather convoluted, but the implications for most large systems are significant: very fast response isn't necessary, and taking time to make good decisions about deployment of FFR is not only desirable from a robustness perspective, but uses FFR resources with finite energy more effectively.

7.2.2 Evaluation of FFR with a Complex Response

In this section, a test case using inertia-based FFR on wind generation (synthetic inertia) is used to illustrate the concept. In the case, the reference disturbance is applied. The sum of the change in power (due to the inertia-based control) for the (equivalent) controlled wind plant is shown in Figure 80.

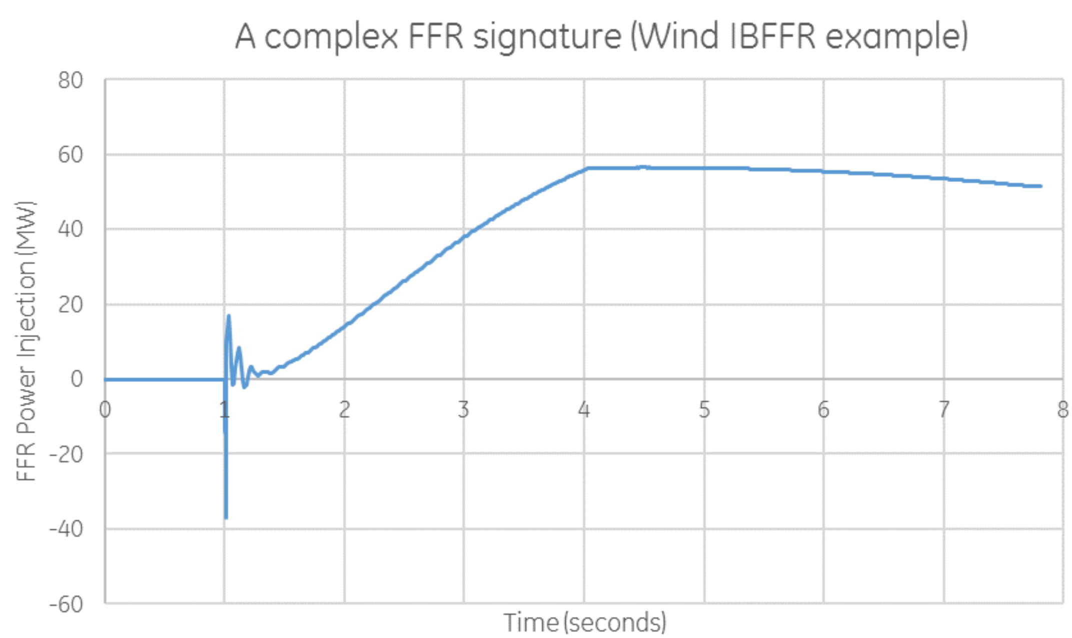

Figure 80 FFR Power with inertia-based FFR from Wind Generation

The arresting energy for each interval of efficacy mapping for this FFR power signature is bucketized and presented in Figure 81.

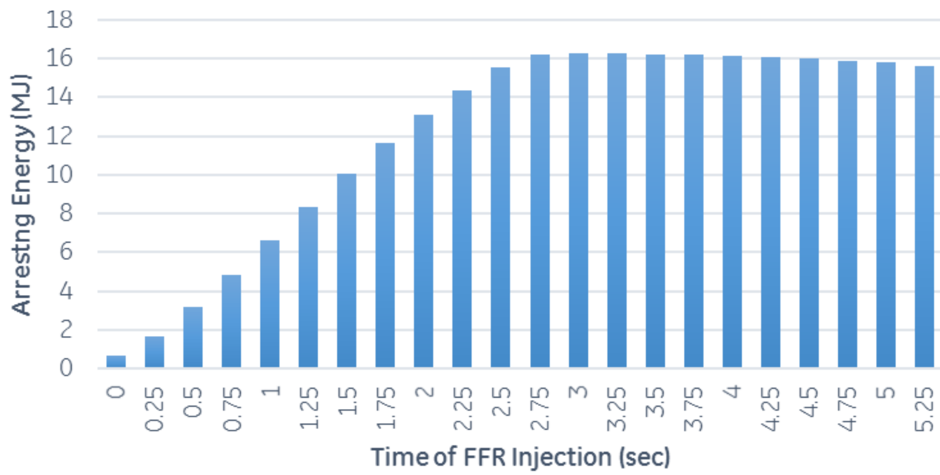

Figure 81 Aggregate Wind IBFFR Arresting Energy

By adding up or convolving the per unit expected benefit (Figure 79) with the amount of energy in each bucket (Figure 81), the impact of each "slice" of FFR energy can be determined, as shown in Figure 82. By adding up the components, the cumulative expected impact on frequency nadir is obtained, as shown Figure 83. So, for this example the *expected* improvement in frequency nadir is 438 mHz.

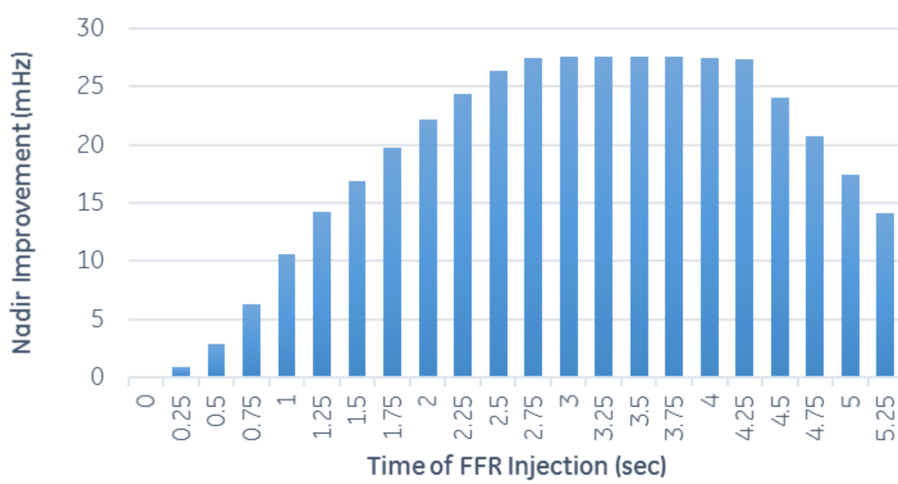

Figure 82 Expected Component Benefit of FFR Energy

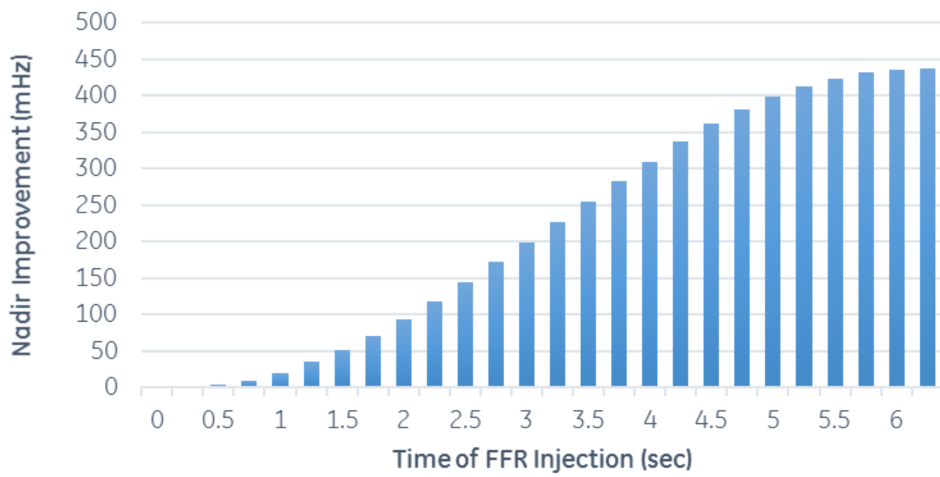

Figure 83 Expected Total Benefit from IBFFR Energy

7.2.3 Confirmation

The frequency response of the time simulation that produced Figure 80 has an actual frequency nadir improvement of 432mHz (0.0086pu), which is essentially exactly as predicted (438mHz) by the linearized method. This provides an encouraging proof of concept, at least in so far as showing that the impact of any FFR signature can be predicted, at least when the expected impact on frequency nadir is relatively small.

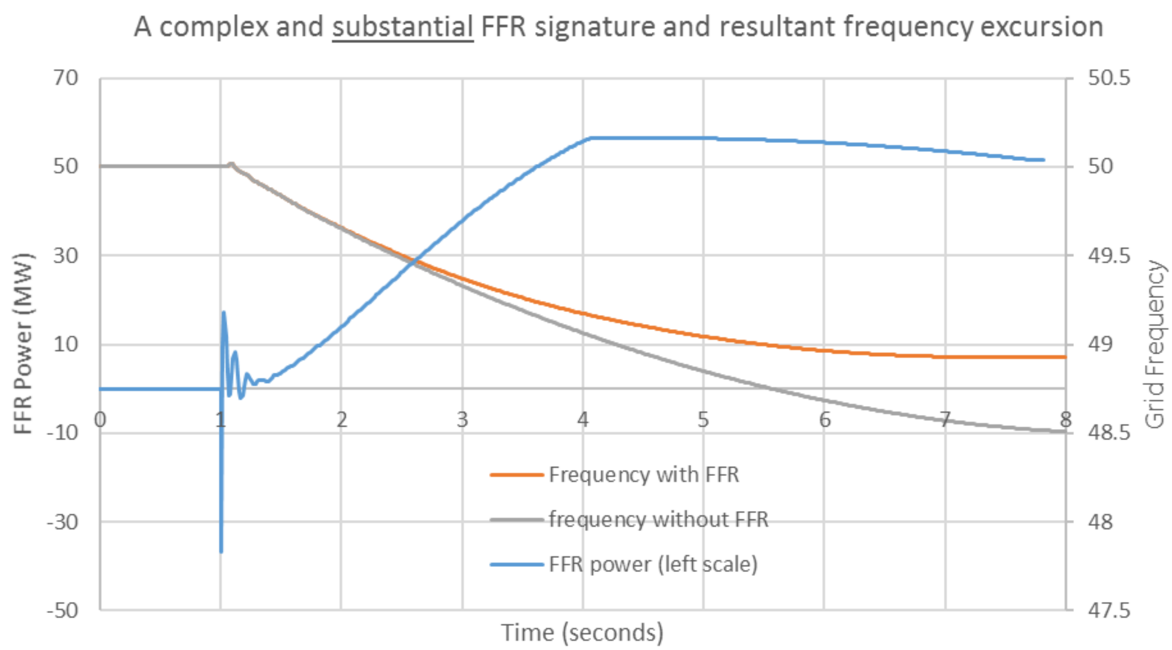

Figure 84 Confirmation of Impact of IBFFR

7.2.4 Non-Linearities and Other Complexities

This method has been tested for much larger individual FFR resources, and as might be expected, linear extrapolation of performance for large signals is imperfect. Tests on large signal (i.e. lots of arresting energy) suggest that the method is rather conservative. That is, this approach tends to somewhat understate the benefit of large FFR arresting energy injections. From a system security perspective, this is highly desirable, as the risk of under-procuring FFR services based on the linear approach appear to be small.

There appear to be promising approaches to improving the method, even for large-signal use. The perturbation method provides useful insight into the impact of different FFR resources on the timing of the nadir, as well as the amplitude. Experiments suggest that this information is useful in refining the estimation of frequency nadir benefit. Further, use of this approach to target other behaviors, such as RoCoF, looks promising. The approach is not a substitute for complete dynamic simulations to validate performance.

7.2.5 Applicability: How High is the RoCoF?

The concept presented here is simple in a world where FFR is the ONLY resource active until the nadir. No coordination with the speed of the PFR is necessary. In other words, FFR does ALL the arresting, stabilizes the system until the PFR activates and takes over. This is the case in systems that have extreme frequency declines (e.g. greater than 1 or 2 Hz/sec…i.e. like South Australia).

But the more common situation is of systems with slower, more typical frequency declines. These systems have historically used only PFR to provide arresting power, but may now find themselves with marginally increasing RoCoFs and marginally short supplies of arresting power. When considering FFR to augment existing (but dwindling) PFR, there is significant value in considering the interaction between the two.

7.2.6 Conclusions of FFR Comparison Method

*A method is proposed here, in which dissimilar sources of fast frequency response can be fa*irly compared. A few key elements emerge from this work:

- Don't forget that the primary objective is to improve the frequency nadir.
- Create an environment where every technology can contribute, and
- Be evaluated on its systemic benefit.
- Faster isn't necessarily better, so don't base requirements or market signals on that assumption.

Next steps: This exercise is illuminating, but there are many complexities to be further considered, including

- How linear are the results?
- How many different conditions need to be considered in order to get robust and workable rules?

- How to keep rules up-to-date?
- Are there unintended consequences?
- How might other objectives be included/rewarded?

Detailed simulations, on a validated system model, are required to answer these concerns.

7.3 A Possible Market Construct

The method outlined above suggests that it is possible to make fair apples-to-apples comparisons of different and complex FFR response. At least when the individual responses being compared are of modest scale relative to system and event size. The question remains, "what to do with this from a practical market or procurement perspective?"

One practical approach would be to use simulation of generic FFR to determine how much is needed to meet performance objectives. This could be as simple as determining how much generic FFR is needed to raise the frequency nadir for the specific operating condition (e.g. as determined by economic security constrained dispatch) for the design event (e.g. interconnector flow) to a level that meets the required frequency standards with some margin (e.g. as set by the rules).

Participants in ancillary service markets could offer amounts and prices for their dissimilar services. The method here would allow the dissimilar services to be compared for their impact on the frequency nadir. The value, and therefore the clearing rank, of the offers would be normalized to cost per unit of improved system performance (i.e. $/ΔFnadir). The market would clear when enough FFR is procured to meet the expected system needs. Of course, a full market isn't needed to realize this benefit. The approach is equally applicable to other longer-term procurement approaches (e.g. tenders or bilateral contracts).

The net result of this approach is that resources with more desirable signatures will get a premium based on their better impact on frequency. But other resources, which might have less ideal signatures, but lower costs, can also participate: More participants equals better economy. Prescriptive rules for specific signatures, that will limit the participation of a wider range of resources, and generally result in less effective use of power and energy ratings of FFR resources, can be avoided.

For the immediate case of South Australia, as noted above, the relationships will tend to become simpler, with resources that provide response fast enough to be useful for the extremes of low inertia having similar value. Unlike the larger interconnected NEM, where the relationship illustrated above will likely hold, South Australia may take a simpler approach to procurement. To wit, there is no need to unnecessarily complicate comparison of speed and shape of response in South Australia because FFR needs to be very fast. In fact, the primary distinction of importance for South Australia will be between discrete, switched resources (i.e. switched blocks of loads) and continuously acting resources (e.g. BESS, IBFFR, etc.). We showed above that at least some response ought to be continuously acting with closed-looped controls. A further distinction may be sustainability of the output; i.e. IBFFR probably needs to be handled separately.

7.4 Procurement Discussion

It is first important to note that physics should drive FFR procurement: maintaining the security of the NEM, especially the South Australian system, is becoming increasingly challenging. This is emphatically a power system dynamics problem that requires good utility engineering practice. Market constructs must be founded on this, and must have sufficient flexibility to adapt to what will inevitably be a rapidly changing situation – both in terms of available technology and details of the grid dynamics.

Provision of FFR from most resources that might be built for the specific purpose of providing that service involves significant investment in physical equipment. For dedicated resources like batteries, flywheels, etc., the capital investment is the dominant cost. The variable cost of providing FFR is usually relatively low. Other resources, which might be able to provide FFR *in addition* to their main function (e.g. IBFFR from Wind, or FFR from temporarily over-driven solar PV), may be able to offer the service after a relatively modest capital investment. Some resources, like batteries, might also incur opportunity costs if they are performing other market functions, such as energy arbitrage[127]. Design of a market for procurement of FFR needs to consider these realities.

Procurement by an annual (or periodic) clearing market (i.e. winning resources are obligated to provide FFR service when called upon, in return for a fixed revenue stream, plus energy net) can be effective. Tendered procurement, in which a call is issued for competitive supply of FFR, can also work. Variable (e.g. Day-Ahead clearing) markets for FFR may face challenges for price formation and stability.

Procurement of FFR resources can follow various paths. The character of the service suggests that the following points apply:

- It will be challenging to get a day-ahead or real-time market to work. Variable (e.g. Day Ahead clearing) markets for FFR will face challenges for price formation and stability.
- Procurement by an annual (or periodic) clearing market could work. In such an arrangement, winning resources are obligated to provide FFR service when called upon, in return for a fixed revenue stream. Payments could also include an energy delivered element.
- Tendered procurement, in which a call is issued for competitive supply of FFR, can also work. This approach might be desirable when the FFR resource is largely separate from other (e.g. energy) markets, and the resource becomes a "network" asset, like transmission or control investment. Such an arrangement would work for bilateral contracts as well.
- The methods outlined in this report for determining the amount of FFR needed to meet performance and economic objectives will provide guidance for establishing how much is needed, and how much can be spent on that procurement (e.g. auction caps, etc.).

[127] Energy storage may participate in energy markets, buying at low prices and selling at high prices. Reserving power rating for FFR could result in lost revenues.

- A part of the market design work should include debating the manner in which loads are involved in FFR. Should participation by loads in triggered response schemes (such as an SPS) be voluntary and paid? Or involuntary and unpaid (as in the existing UFLS scheme)?

The requirements for FFR will continue to evolve in all parts of the NEM, including South Australia. Processes will be necessary to keep the technical evaluation of the frequency response requirements current. An annual review, at least, is probably required.

7.5 Plant Connection Mandates

Some of the technologies discussed in this report, most notably IBFFR from wind generation and some types of load control, involve relatively inexpensive improvement to equipment. Several system operators around the world have implemented requirements (i.e. grid codes) that mandate that wind plants have IBFFR functions to interconnect. This was discussed in Section 3.2.5.

8 RECOMMENDATIONS FOR FUTURE ANALYTICAL WORK

The section includes recommendation for future analytical work, primarily simulations, that will aid in making good decisions regarding FFR services.

8.1 Recommendations for Detailed Dynamic Analysis

8.1.1 General

- Illustrative simulation results used in this work are not a substitute for detailed, full scale dynamic analysis.
- Most analysis can be performed on fundamental frequency positive sequence tools (e.g. PSS/e)
- At very low inertia and low short circuit current levels, at least some analysis to assure control stability (e.g. EMT studies) is warranted.
- Broadly, we recommend that the analysis included in the *illustration* presented in Section 6 be performed on the "real" system. A few specific analyses will be of considerable value, including the analyses that produced Figure 72 and Figure 73.

8.1.2 Class of Analysis/Cases

Combined Fault and Trip Events.

- Events which cause islanding of South Australia are likely to involve transmission faults.
- Dynamics of all systems elements are complicated by faults.
- No types of generation can inject active power into deeply depressed voltages; attempts to do so tend to exacerbate transient stability and transient voltage collapse problems.
- Realistic fault and trip to island events and realistic loss of generation (e.g. in South Australia) events need to be included in the mitigation design process.
- Issues of maintaining control stability for such extreme conditions are much more than just high RoCoF.
- Stability considerations are likely to dominate functional and equipment specification of any remediation technologies.

Dynamic/Control Stability.

- As noted, issues of maintaining control stability for such extreme conditions are much more than just high RoCoF.
- Dynamic stability must be assured. Therefore:
    - Frequency domain analysis is required.
    - Eigenvalues must be stable before, and at projected quasi-static conditions following disturbances.
    - Linearization of non-conventional technologies is required, and can be a non-trivial challenge.

- Modal participation and eigenvalue migration may be surprising under such substantially different conditions.

- Simulations that will support the design of the recommended SPS are needed. Specifically, design of the SPS will require that:
  - The necessary speed of load tripping be established,
  - All events that can lead to islanding be identified,
  - Better understanding and verification of the requirement for exactness of "equal" as a function of inertia and import be developed.
- A centrally armed autonomous SPS, using local triggering (section 6.6) will need simulations to establish whether robust performance can be obtained under low inertia conditions for the possible spectrum of separation events.

8.1.3 Recommendation for Modeling for Detailed Dynamic Analysis

Good model fidelity is essential. Creating and maintaining high fidelity dynamic models is (and must be) an ongoing process. It requires significant effort. Validation should be made:

- Against observed events; against validated/confirmed field settings (e.g. for UFLS relays; generator protection).
- For staged, field tests (e.g. like NERC model validation testing[128]).
- Of load response (this is difficult; a program of monitoring should be considered).
- Of dynamic performance of all resources providing FCAS, with attention to assessing how those resources respond to frequency error.

Calibration analysis

- Illustrative analysis included several parametric variations that could be valuable if performed on the full South Australia system model.
- If RoCoF measurement and event detection is used as a main method for triggering FFR (not our recommendation), frequency from real events should be examined (at multiple locations) to determine the size of the sampling window to manage the tradeoff between accuracy in event detection and time to respond.

8.1.4 Recommendations for other technical analysis and simualtions

Short circuit screening

- Evaluation of grid strength and minimum short circuit strength for future conditions is needed.
- Screening for short circuit levels below those acceptable to equipment manufacturers or below current technology tolerances are needed. These might point to the need for EMT analysis

128 http://www.nerc.com/pa/Stand/Project%20200709%20%20Generator%20Verification%20%20PRC0241/MOD-027-1_clean_2012Sept11.pdf

ElectroMagnetic Transients analysis

- Depending on screening results, as future very low short circuit strength scenarios become defined, EMT analysis of the behavior of specific inverter-based resources, (including, but not limited to new FFR technologies) will be required.
- Analysis for risk of poor control performance due to weak grid may be required, but after basic fundamental frequency and frequency domain work has established credible alternatives.
- Detailed modeling of MurrayLink in a suitable platform (e.g. PSCAD) to determine the possibility of using the project as a source of FFR should be considered.

8.2 Recommendations for Detailed Economic Analysis

8.2.1 General: Production Simulation

Evaluation of the economic trade-offs between various alternatives done based on high fidelity production simulations will give the most solid foundation for decision making.

For South Australia, production simulations that:

- Use real (or realistic) hourly wind and solar production
- Correctly constrain operation of the AC interconnector subject to the inertia and PFR provided by synchronous generation in South Australia (per discussion in Section 6.3.5)
- Constrain operation subject to the allowable level of involuntary load-shedding
- Correctly capture the contribution of various alternatives (e.g. new supplies of FFR and improved PFR) in partially or completely relieving constraints

This type of simulation will provide reference values for the impact on the annual variable cost of energy, impact on carbon and other emissions. This type of analysis is essential to determining the relative value of different strategies. Strategies could include addition of capital equipment, more conservative operation, relaxed reliability or performance expectations. It is not meaningful to only ask the capital cost of adding different widgets!

Simulations of these types normally assume that economic responses will be competitive, and that markets has sufficient depth to avoid market power. Additional investigation of the robustness of price formation for the operating conditions calculated by the recommended production simulations may be warranted.

8.3 Zero Synchronous Generation Future

As noted in Section 4.5, this a future condition that needs to be considered. Technology and solutions are far outside the narrower scope of this FFR investigation, but the state of the art is changing very rapidly. South Australia should, for now, proceed with the understand that there should be no expectation *today* that it can survive, islanded from the rest of the NEM, without at least some synchronous equipment in operation. To make zero synchronous generation operation possible in the near future, investment will be required, some of which may be in technology that is not commercially available today.

9 REFERENCES

In addition to footnotes throughout the document, additional references are provided here.

9.1 RoCoF and Frequency Response References

DNV-GL, *RoCoF Alternative Solutions Technology Assessment – Final Report Phase 1,* Report No. 16011111, Rev. 005, 17 August 2015.

EPRI/DOE Handbook 2003

DOE/EPRI 2013 *Electricity Storage Handbook in Collaboration with NRECA*, SAND2013-5131, July 2013.

Application of rate of change of frequency constraints for high wind penetration scenarios in a small power system, George Ivkovic, Power Engineering Conference (AUPEC), 2013 Australasian Universities, Year: 2013, Pages: 1 - 5, DOI: 10.1109/AUPEC.2013.6725374

Network Control System Protection Scheme: Design enhancement, Tien Ho, Power Engineering Conference (AUPEC), 2013 Australasian Universities, Year: 2013, Pages: 1 - 6, DOI: 10.1109/AUPEC.2013.6725436

Advances in Wide Area Monitoring and Control to address Emerging Requirements related to Inertia, Stability and Power Transfer in the GB Power System, D.H. Wilson et al, CIGRE Paper C2-208, Paris, August 2016.

Smarter Network Storage: SDRC 9.7 Successful Demonstrations of Storage Value Streams, P. Papadopoulos, A. Laguna-Estopier, and I. Cooper, UK Power Networks; March 2016; Pages 1-71. http://innovation.ukpowernetworks.co.uk/innovation/en/Projects/tier-2-projects/Smarter-Network-Storage-(SNS)/Project-Documents/SDRC+9.7+Successful+Demonstrations+of+Storage+Value+Streams+LoRes+v1.pdf

J. Eto, et al. "Use of Frequency Response Metrics to Assess the Planning and Operating Requirements for Reliable Integration of Variable Renewable Generation", LBNL-4142E, Dec. 2010. http://www.ferc.gov/industries/electric/indus-act/reliability/frequencyresponsemetrics-report.pdf

N. Miller, et al, "California ISO (CAISO) Frequency Response Study," Nov. 9, 2011. https://www.caiso.com/Documents/Report-FrequencyResponseStudy.pdf

N. Miller, et al, "Eastern Frequency Response Study," NREL//SR-5500-58077, May 2013. http://www.nrel.gov/docs/fy13osti/58077.pdf

N. Miller, et al, "Western Wind and Solar Integration Study Phase 3 – Frequency Response and Transient Stability", NREL/SR-5D00-62906, December, 2014. http://www.nrel.gov/docs/fy15osti/62906.pdf

9.2 Frequency-Input Controls for Damping

Robust damping controls for large power systems, J. F. Hauer, IEEE Control Systems Magazine, Year: 1989, Volume: 9, Issue: 1, Pages: 12 - 18, DOI: 10.1109/37.16744

Robustness issues in stability control of large electric power systems, J. F. Hauer, Decision and Control, 1993., Proceedings of the 32nd IEEE Conference on, Year: 1993, Pages: 2329 - 2334 vol.3, DOI: 10.1109/CDC.1993.325613

9.3 HVDC

Implementations and experiences of wide-area HVDC damping control in China Southern Power Grid, Chao Lu; Xiaochen Wu; Jingtao Wu; Peng Li; Yingduo Han; Licheng Li, 2012 IEEE Power and Energy Society General Meeting, Year: 2012, Pages: 1 - 7, DOI: 10.1109/PESGM.2012.6345363

Description of the new RAS to be installed in the power system of Uruguay, D. S. Beledo; D. Bonjour; F. Sanchez; N. Yedrzejewski, Innovative Smart Grid Technologies Latin America (ISGT LATAM), 2015 IEEE PES, Year: 2015, Pages: 188 - 193, DOI: 10.1109/ISGT-LA.2015.7381151

Designing the hardware and software of the Pacific HVDC Intertie Remedial Action Scheme using a programmable controller, M. R. Bhuiyan; A. E. Taylor; N. C. Kezman, Transmission and Distribution Conference, 1991., Proceedings of the 1991 IEEE Power Engineering Society, Year: 1991, Pages: 790 - 796, DOI: 10.1109/TDC.1991.169594

Latest control and protection innovations applied to the Basslink HVDC interconnector, M. Davies; A. Kolz; M. Kuhn; D. Monkhouse; J. Strauss, AC and DC Power Transmission, 2006. ACDC 2006. The 8th IEE International Conference on, Year: 2006, Pages: 30 - 35, DOI: 10.1049/cp:20060007

The Intermountain Power Project 1600 MW HVDC transmission system, C. T. Wu; P. R. Shockley; L. Engstrom, IEEE Transactions on Power Delivery, Year: 1988, Volume: 3, Issue: 3, Pages: 1249 - 1256, DOI: 10.1109/61.193910

Benefits of HVDC for reducing the risk of cascading outages and large blackouts in AC/DC hybrid grid, M. Benasla; T. Allaoui; M. Brahami; A. Boudali, Control, Engineering & Information Technology (CEIT), 2015 3rd International Conference on, Year: 2015, Pages: 1 - 6, DOI: 10.1109/CEIT.2015.7233170

Thyristor valve for the 12-pulse converter for the Champa-Kurukshetra HVDC transmission scheme, Mohammad Hassan Jodeyri; Andrzej Dzus, 2013 IEEE Innovative Smart Grid Technologies-Asia (ISGT Asia), Year: 2013, Pages: 1 - 6, DOI: 10.1109/ISGT-Asia.2013.6698790

MurrayLink, The Longest Underground HVDC Cable in the World, I. Mattsson et al, CIGRE Session B4-103, 2004

Power Modulation of Sidney HVDC Scheme, part 1, RAS Control Concept, Realization, and Field Tests, RK Johnson et al, IEEE Transactions on Power Delivery, Vol.4, No.4, October 1989.