MARKET REFORM

OPERATIONAL SECURITY MECHANISM

Modelling Findings Report

FINAL 10 MARCH 2023

AEMO
AUSTRALIAN ENERGY MARKET OPERATOR

NOTICES

Disclaimer

This report has been prepared by Market Reform at the request of AEMO. The report explores potential alternative OSM design and formulation options and provides insight into their potential performance and impacts for different hypothetical cases. The report is intended to provide information to inform future discussion around OSM but is not intended to provide predictions of any actual future market outcomes. We do not accept any liability if this report is used for an alternative purpose from which it is intended, nor to any third party in respect of this report.

EXECUTIVE SUMMARY

Overview

The National Electricity Market (NEM) generation mix consists of a significant and growing proportion of inverter-based resources (IBR), i.e., wind and solar PV. IBR do not interact with the power system in the same way as traditional synchronous resources (coal, gas and hydro). As these synchronous resources are displaced, it can become more challenging for system security constraints to be met, with the existing market design and spot market prices not necessarily incentivising a secure combination of resources to be available.

To ensure that the power system can be operated securely as IBR capacity continues to grow, and to ensure that provision of required security services is incentivised and valued, a so-called Operational Security Mechanism (OSM) is being explored by the Australian Energy Market Commission (AEMC). The OSM would run ahead of the real-time dispatch and commit units where required to meet minimum security constraints and/or to relieve constraints on IBR if this enhances social welfare.

Market Reform has been engaged by the Australian Energy Market Operator (AEMO) to formulate and develop a toy OSM model and formulation. This model is capable of exploring various scenarios, including how the OSM interacts with the NEM pre-dispatch (PD) and participant decision-making, as well as the impact of various scheduling parameters (e.g., simulation horizon and resolution). The model has been intended to inform – but not directly be used in - the final OSM solution.

Model and data

The prototype models both the OSM itself, as well as pre-dispatch processes before and after the OSM is run. It represents the five regions in the NEM, and includes all scheduled and semi-scheduled units currently operating, as based on the 2022 Integrated System Plan dataset. Demand, IBR forecasts and energy and FCAS offers were obtained for selected days for the period of late 2021 - late 2022, with days selected for conditions that are likely to require OSM commitments. A simplified transmission model was used.

Various types of OSM security constraints were represented. The focus was on secure configurations in Victoria, and particularly South Australia. In the latter region, different configurations allow for different constraints on wind and solar output, so a may be worth incurring additional OSM commitment costs if this is more than offset by the value of low-cost energy. Minimum numbers of synchronous units were also required in each region. Finally, in certain case studies, South Australia was assumed to be islanded, and a local inertia constraint was applied.

The formulation also tested the ability to schedule gas or hydro units in synchronous condenser mode (in which they were assumed to have a lower security contribution, but with avoided fuel/water costs). Additionally, assumptions were made about batteries being able to operate in a mode to provide inertia. These assumptions were only for testing and validation purposes, and not intended to represent the current or anticipated engineering knowledge or indeed future commercial uptake of these technologies.

Case studies

A number of case studies were examined to test the model and explore design questions. Most notably:

- The OSM was observed to make commitments both for minimum security reasons, or to incur additional commitment costs if this was more than offset by improvements in the gains from trade.
- Units were committed in synchronous condenser mode to meet a separate inertia constraint when South Australia was assumed to be islanded.
- It was not observed that OSM units with a minimum generation constraint were unable to be scheduled at this level in the PD run after the OSM, if they offer at the price floor, except in a case that was specifically set up to produce this outcome by using extremely low offer prices for wind.
- In a simple worked example, disorderly bidding from IBR units was observed to result in commitments of OSM units that otherwise would not improve the gains from trade were offers to reflect true costs.

Scheduling parameters

One of the OSM design questions relates to the choice of the so-called scheduling parameters. These parameters – which are shown conceptually in Figure 1 – can have impacts such as on the economic efficiency of the OSM schedule, the OSM solution time and the volatility of OSM decisions (e.g., due to changes in inputs like renewable forecasts). They may also have other impacts – such as on complexity of structuring competitive OSM offers for participants – that cannot be easily tested in the toy model.

Figure 1 – Conceptual view of scheduling parameters

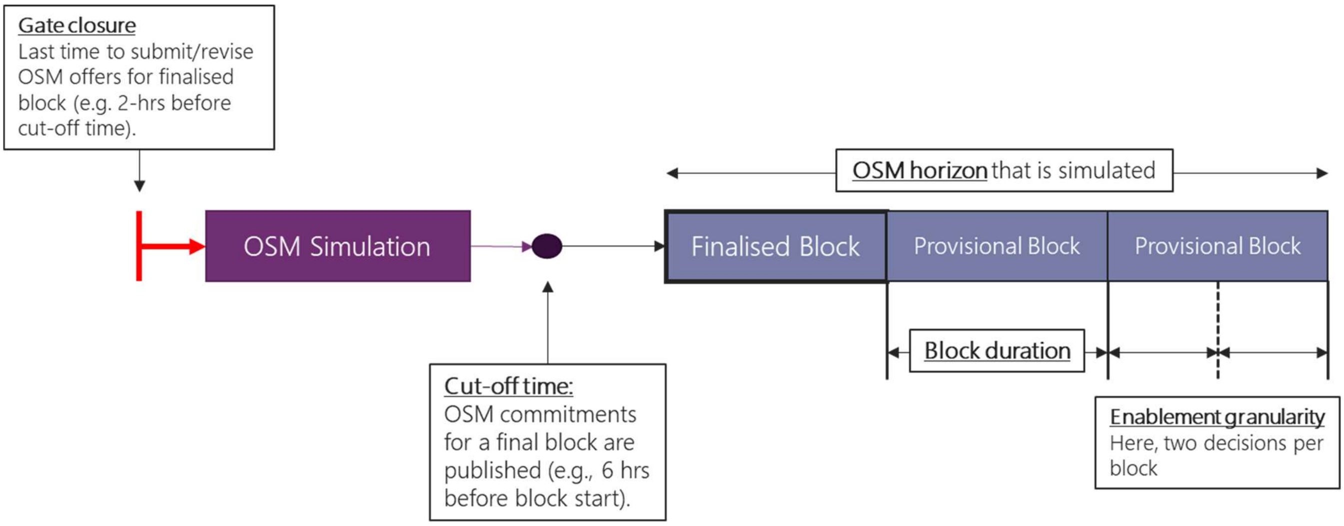

Table 1 shows the scheduling parameter conditions that were tested, as well as high-level observations. Overall, multiple combinations of these scheduling parameters were found to meet criteria of a) avoiding significant loss of economic efficiency b) producing reasonable run times[1] and c) without creating overly variable or unpredictable outcomes. The findings suggest that a block size of 4 or 8 hours, with resolution of 2-hours would produce a reasonable trade-off of the various considerations, but is not the only workable combination. However, it is recommended that the option of optimising over only a single block not be progressed.

[1] Note that solution times for the toy model may not be indicative of that for a real implementation of the OSM.

Table 1 - Scheduling parameters for testing

PARAMETER	TESTED CONDITIONS	OBSERVATIONS
Block size	2, 4 and 8 hours,	Optimising over a longer duration means decisions are made with more uncertainty, but the difference in solution quality with 4-hour or 8-hour block duration is small, and both of these options appear workable.
Granularity of OSM enablements	30-min, 2, 4 and 8 hours	There were only small benefits from using a duration less than 2-hours, and this may partially be because many of the core units in the OSM configurations have minimum on times for 4 or more hours. Additionally, many of the security gaps were for periods longer than 2-hours, so there is little benefit in making decisions that change at this granularity.
Treatment of horizon	Either only the first block is simulated, or all blocks in a rolling 24-hours are simulated (e.g., 3 eight-hour, or 6 four-hour blocks)	The single block approach was observed to result in less optimal decisions, as there is no information about the influence of decisions made now on decisions that can or must be made in the future. This is particularly the case for a short block duration. The single block approach is not recommended.
Cut-off time	1, 2 and 4 hours	Making decisions further ahead of block commencement increases uncertainty in wind, solar and demand. This produced worse outcomes, but not prohibitively so. Making decisions further ahead may also make participation easier for slow-start units, and give more operator and participant certainty.

Summary of findings and recommendations

High-level findings and recommendations are as follows:

- Subject to points raised below, the OSM formulation itself, and its integration within the PD and other scheduling processes was found workable end-to-end. The formulation appropriately made least-cost decisions to meet minimum security requirements, or incurred additional costs where these allowed for overall improvements in the gains from trade. The formulation is able to consider both conventional unit commitment decisions, as well as more novel units such as dual-mode units and batteries providing inertia.

- The AEMC's draft determination envisages that the algorithm should avoid OSM commitments that are made only to improve the gains from trade through the OSM unit's energy, rather than by contributing to meeting or relieving security constraints. This is termed an energy-only commitment and might occur where a unit's OSM offers are less than the marginal energy price, and – for whatever reason – that unit has chosen not to self-commit. A two-step solution to this problem has been proposed, though not implemented or tested in the model. Further work is recommended to a) test this solution and limitations, and b) more precisely define the scenarios in which it is or is not acceptable to commit a unit through the OSM.

- A related issue is that disorderly bidding could result in OSM commitments that would not be economic under bids that reveal true preferences to generate or not. In other words, disorderly bidding could be used to make OSM commitments that do not make the system as a whole better off, but do make certain participants better off. This can arise because of the NEM's regional pricing structure (which is known to produce disorderly bidding incentives even without the OSM), but potentially also because the costs of an OSM commitment may not necessarily be recovered from those that benefit from it. While this was observed as a possible outcome, it is not necessarily concluded that it would occur commonly in practice.

- With regards to scheduling parameters, it is recommended that the approach be to incorporate at minimum a 24-hour horizon rather than a single block scheduling approach, as this avoids end-effects that occur with a shorter horizon. Beyond this:
  - Results indicate there is only small additional benefit in using a commitment granularity less than 2-hours, though there may be little downside in doing so except for longer solution times. Commitment granularity longer than two hours results in approximately 0.8% – 1.5% lower improvements in the gains from trade.
  - Results with either 4-hour or 8-hour block durations are found to give similar improvements in the gains from trade.
  - There are multiple combinations of parameters that are workable within the prototype, and other factors – particularly impacts to operator and participant decision-making - should be considered
- Solve times for our prototype are typically approximately 3-minutes, using a standard desktop PC and a commercial optimisation solver, but with some runs requiring up to 14 minutes, and with this variation primarily being impacted by the inputs (demand, wind and solar forecasts etc) themselves, rather than the scheduling parameters. It is cautioned that solve times may differ in a production version.
- Consideration should be given to the extent to which the OSM scheduling algorithm should represent the full range of generic transmission and other constraints that are included in NEMDE and pre-dispatch models. These constraints would introduce new variables that may increase solution time, yet may not all be required for a reasonable solution.

CONTENTS

TABLES

FIGURES

ABBREVIATIONS

ABBREVIATION	TERM
AEMC	Australian Energy Market Commission
AEMO	Australian Energy Market Operator
CNQ	Central North Queensland
CNSW	Central New South Wales
COAG	Council of Australian Governments
ESB	Energy Security Board
FCAS	Frequency Control Ancillary Services
FFR	Fast Frequency Response
ISP	Integrated System Plan
LP	Linear Program
MILP	Mixed Integer Linear Program
NEM	National Electricity Market
NEMDE	National Electricity Market Dispatch Engine
NNSW	Northern New South Wales
OSM	Operational Security Mechanism
PD	Pre-Dispatch
PV	Photovoltaics
SC	Synchronous Condenser
SG	Synchronous Generator
SNSW	Southern New South Wales
SNW	Sydney, Newcastle and Wollongong
SQ	South Queensland
UC	Unit Commitment

1 INTRODUCTION

The National Electricity Market (NEM) generation mix consists of a significant proportion of inverter-based resources (IBR), i.e., wind and solar photovoltaics (PV), with this proportion expected to grow in coming years and decades. IBR do not interact with the power system in the same way as traditional synchronous resources (such as coal, gas and hydro resources). As synchronous resources are displaced, it has become more challenging for system security constraints to be met. These constraints are an active area of research, but may relate to system strength, inertia, and other considerations.

At times, energy prices alone may not provide a sufficient incentive for resources to be online such that security constraints are met. At other times, the system may be secure, but constraints must be placed on how much generation from IBR can be dispatched. As such generation is typically low cost, this may then have economic consequences.

To ensure that the power system can be operated securely as IBR capacity continues to grow, and to ensure that the security contribution of relevant resources is valued, a so-called Operational Security Mechanism (OSM) is being explored. The OSM would run ahead of the real-time dispatch in the NEM Dispatch Engine (NEMDE), on the same/similar timeframe (or in tandem with) as the Pre-Dispatch (PD). Registered units would be offered into the OSM with prices for these units to be committed. The OSM would then commit some of these units where required to meet minimum security constraints, and/or to relieve constraints on IBR where this improves the so-called gains from trade (gains from trade is defined below, but loosely is the difference between the value of serving demand, less the cost of operating resources to meet that demand). A mathematical optimisation program would be used to make these commitment decisions.

The introduction of the OSM to the NEM is being progressed in a rule change process by the Australian Energy Market Commission (AEMC). Currently, a draft determination has been published (21st September 2022), with a final determination due on the 27th July 2023.

Market Reform has been engaged by the Australian Energy Market Operator (AEMO) to formulate and develop a prototype OSM model/formulation. This model is to be used to explore various case studies, including how the OSM interacts with the NEM PD and participant decision-making, as well as the impact of various scheduling parameters (e.g., time horizon and resolution). This work is based on the OSM draft determination, and is intended to understand the implications and feasibility of different approaches, inform design trade-offs, and also to help plan for eventual implementation.

This document reports on findings in developing the OSM prototype. It is laid out as follows:

- Section 2 introduces the OSM and the study aims.
- Section 3 describes the modelling approach and assumptions, etc.
- Section 4 summarises the formulation, with additional detail in the appendices and accompanying formulation prototype user guide.
- Section 5 sets out the data used in the modelling.
- Section 6 presents modelling case studies and exploration of design questions.
- Section 7 assesses the various scheduling parameters (block duration, granularity of enablements, etc)
- Section 8 summarises the findings and makes recommendations.

Finally, Appendix A sets out the detailed mathematical description used for the OSM optimisation formulation.

2 OSM OVERVIEW AND AIMS OF THIS STUDY

2.1 OSM overview

The AEMC's proposed OSM would be a process that is run in the hours/days leading up to real-time dispatch, to be used to ensure that a combination of resources is available with the required attributes and capabilities to meet system security constraints. As well as ensuring these resources are available, it provides a means to value security contributions and remunerate costs, and therefore can be viewed as a market for security. This section provides a brief description of the features of the envisaged OSM that are relevant to this study.

For the power system to be secure, various requirements relating to – for example - voltage, frequency and line flows must be met. As inverter-based capacity in the NEM increases, requirements to meet these security constraints have become an evolving area of research.

Some ancillary services are already procured in the real-time spot market, namely the various frequency control ancillary services (FCAS). The OSM would not procure these spot-traded services. Rather, it would optimise resources to meet other security requirements that may be less well-suited to real-time trade. These requirements could include ensuring that discrete configurations of units are available that are known to place the system in a secure state, or be separate service requirements such as for power system inertia.

The OSM would run in tandem with the existing pre-dispatch (PD) process, which is a non-binding process that provides a simulated view of how the market would clear and units would be dispatched using the most recent demand forecasts, constraints, unit offers, and so on. The PD is currently run every half-hour, and allows participants to understand potential pricing and dispatch outcomes, and AEMO to assess security and reliability considerations.

Units that can contribute to security needs would need to be accredited for participation in the OSM. Most OSM units might be synchronous generating units but other unit types could potentially participate. For example, some generating units may be able to contribute to security in synchronous condensers, or batteries may be capable of providing inertia or other services. For these units, the OSM would be able to choose which mode it schedules the unit in, subject to some constraints and OSM bids to operate in each mode by the resource operator.

OSM units would submit offers to be scheduled in the OSM. These would differ from PD offers:

- PD offers are up to ten price-quantity pairs for energy and for each FCAS.
- OSM offers are formed of a running cost ($/MWh at minimum enablement) and a start-up cost ($/start).

On certain days, some OSM units may prefer not to be scheduled via the OSM – for example if they expect the spot market energy price would be sufficient to recover their costs. Therefore, they would submit a self-commitment flag for each period, indicating whether they wish to be considered in the OSM (if the flag is false).

The OSM would run several times in the lead-up to real-time, providing additional information to the PD. Initial runs of the OSM would be non-binding – e.g., a cleared unit is not yet required to run in real-time, and the OSM is only indicating it is expected to be committed (all else being equal). Nonetheless, the unit's commitment would then be expected to be reflected in the subsequent PD runs (either via updated offers or constraints). At some point ahead of real-time dispatch, the last OSM is run for a particular period, and any commitments for that period become binding. The chosen units are required to operate in accordance with those commitments, and would be paid according to their OSM offers (OSM and spot market settlement calculations are described in detail in Section 3.4).

Figure 2 presents a conceptual view of how the OSM and PD interact in the lead-up to dispatch in real-time, and settlement against the final schedules.

Figure 2 - Sequence of PD and OSM runs, with settlement occurring once schedules are final

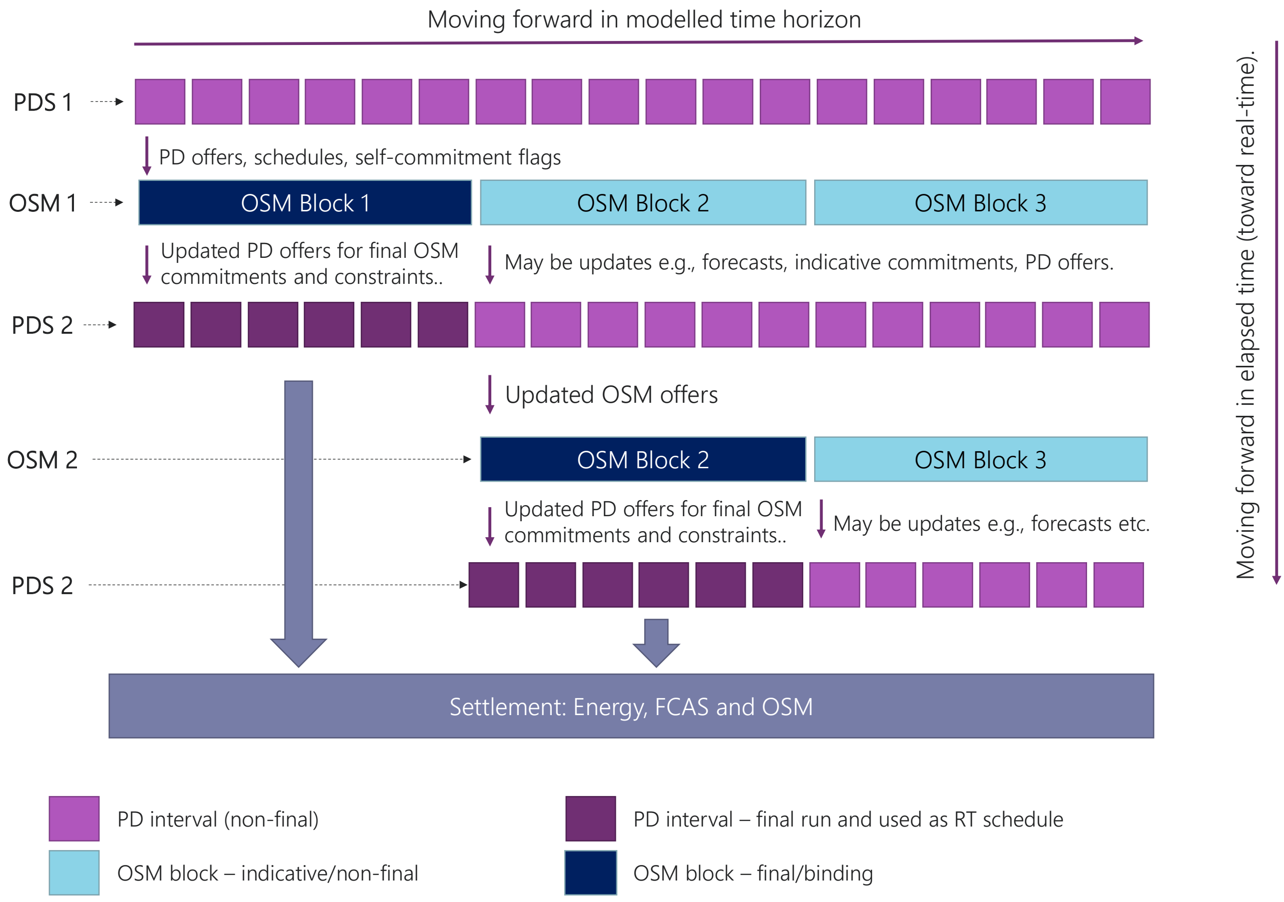

2.2 Study objectives

This work sets out to build on earlier work by formulating and building a working toy model of the OSM, and using this model to explore the performance and behaviour of the OSM. This also includes exploring how the OSM would interact with the PD (both from the point of view of scheduling OSM units in a subsequent pre-dispatch, but also how the PD might be used by participants to make decisions about setting their self-commitment flags).

A number of case studies and design questions were defined to be explored with the model, for example:

- Should the OSM mechanism use bidding at the price floor in PDS as an approach to ensure resources cleared in the OSM are scheduled in dispatch?
- What are the mechanics of the OSM relieving Inverter Based Resources (IBR) constraints for net market benefit?
- What are the mechanics of the OSM unbundling system services from configurations?
- How does the OSM treat units that can operate in multiple modes (e.g., fast-start units that can operate in synchronous condenser mode)?

The study is also concerned with assessing the relative performance of the prototype under different scheduling parameters – for example the granularity of time blocks in the OSM, and the number of blocks that are solved in an OSM run.

In the course of this work, observations and findings have and are have been discussed with the market bodies and interested stakeholders, allowing these observations to be considered with respect to both the OSM design and implementation.

Note that this body of work is not intended to make predictions about the future, such as how many units, or which type of units would use the OSM. Rather, it is intended to develop a formulation that can support the policy objectives with a broad range of possible technology types. As an example, to test functionality relating to choosing whether a unit should be operated in synchronous generation or synchronous condenser mode, it is assumed that some participants have enabled their units with this capability. This assumption is made for the purposes of testing the model only.

3 MODELLING APPROACH

To explore the OSM formulation and performance, a toy model of the OSM, and the broader NEM scheduling processes was developed. This section discusses the modelling approach, while the following Section 4 presents a detailed mathematical OSM formulation.

3.1 Model overview

The model represents each of the five regions in the NEM (Queensland, New South Wales, Victoria, South Australia and Tasmania), and includes the scheduled and semi-scheduled generating units in the NEM (as in the 2022 Integrated System Plan database). A small number of hypothetical units were also used to explore possible scenarios that may occur in the future.

The model is formulated as a series of Linear Programs (LPs) or Mixed Integer Linear Programs (MIP). These simulate the OSM, as well as the PD outcomes both before and after the OSM has been run. However, it was also necessary to include other model runs, to ensure feasibility of inputs between – for example – the first PD run, and the subsequent OSM run. Therefore, five distinct runs – referred to as phases - are used in the modelling. The third phase represents the OSM itself, and is the primary area of interest. Other phases represent the PD, as well as simulating how participants might expect to behave:

- An initial PD run that clears the market using PD offers and does not include OSM security constraints.
- An adjustment to the previous run to ensure PD schedules respect UC constraints. The output of phase 1 and 2 is the self-commitment schedule, and context of outcomes without the OSM
- The OSM run, which includes PD constraints but also OSM security constraints and unit constraints e.g., unit commitment and dual mode operation. This phase commits non-self-scheduled OSM-units as needed to meet or relieve security constraints.
- Phase 4 is similar to Phase 2, being a PD run with UC constraints, but with the OSM schedules reflected (either via constraints or modifications to PD offer data, e.g., submission of a unit's minimum generation quantity at the price floor).
- Phase 5 repeats Phase 4 but without the UC constraints, representing the updated PD schedules and prices.

The model is capable of reporting on a wide range of outputs from each phase, with primary quantities of interest being:

- Net social welfare at each phase
- Schedules in the PD and OSM
- Prices for energy, and costs of OSM commitments.
- Security violations in the PD
- Constraints on IBR
- Settlement outcomes.

The model was developed using the General Algebraic Modelling System (GAMS) language, with Microsoft Access used for the input and output data storage and reporting. MS-Excel was used for analysis and visualisation of results.

Two solvers were used. GAMS – XPRESS MIP is a commercial solver which was used for most model runs and all performance tests. Another solver, CBC (COIN-OR Branch and Cut), which is an open-source MILP solver that comes with GAMS, was used for additional analysis or model refinements.

- With XPRESS, solve times for the OSM phase (on a desktop PC) have been observed up to six minutes.

- With the less performant CBC solver, solve times for the OSM phase were – at times - long enough to be impractical (>1 hour). However, with smaller numbers of units participating in the OSM, solve times were a few minutes.

Note that solution times for other phases are substantially quicker than the OSM phase. Further discussion on solve times (including the influence of the scheduling parameters) is provided in Section 6.

3.1.1 OSM security constraints

Three types of security constraints are required to be met by the OSM prototype:

- **Configuration constraints:** A requirement that a region has a valid combination of units committed, such that it is in a valid configuration. These configurations have currently been determined for the South Australian and Victorian regions only, and so only apply in these regions. As an example, valid configurations for South Australia include two Torrens Island B units, 3 Dry Creek units, and 1 Mintaro unit (SA_78), or 2 synchronous condensers (at Davenport or Robertson), 2 Torrens Island B units, and 1 Pelican Point unit.
  - There are separate configurations depending on whether South Australia is deemed to be in an islanded state, or not.
  - Configurations can be associated with other constraints. In South Australia, different configurations apply different upper limits on generation from South Australian IBR (wind and solar). In Victoria, there can be import limits (of 850 MW or 800 MW combined from Tasmania and NSW) associated with six configurations. However, most Victorian configurations have no transmission limit, so the South Australian configurations are the main area of interest.

- **Minimum synchronous units constraint:** A requirement on minimum number of particular synchronous units. These limits apply in all regions, but are particularly used in place of a configuration constraint for Queensland, NSW and Tasmania. Regional requirements are shown in Table 2.

- **Inertia constraint:** A requirement to meet a minimum level of power system inertia in a particular region. Requirements are to be determined in specific scenarios, but may be informed by the secure operating level of inertia.[2] Synchronous generators and synchronous condensers provide inertia when committed. In addition, some batteries will be modelled as being able to provide inertia when they have sufficient head and foot room to increase/decrease output in response to a generation-load imbalance.

By default, only the first of these two constraints (configurations and minimum synchronous units) are included in the simulations. The inertia constraint will be included in certain scenarios only, e.g., when a certain region is taken to be in an islanded state, such that all inertia must be supplied locally.

Table 2 - Minimum synchronous units

REGION	MINIMUM SYNCHRONOUS UNITS
New South Wales	7
Queensland	11
South Australia	2
Tasmania	3

[2] Australian Energy Market Operator - 2022 Inertia Report, December 2022

REGION	MINIMUM SYNCHRONOUS UNITS
Victoria	5

While the intent is for the model to be a representation of the real NEM, the objective is to test performance of the OSM under various scenarios. Therefore, in some simulations, requirements or other parameters may be varied away from typical real-world values to examine how the OSM would perform in a hypothetical future world. This also includes making assumptions about how units may be capable of contributing to security requirements in the future – for example dual mode units or provision of synthetic inertia.

3.1.2 Simplifications and assumptions

The modelling approach includes a number of simplifications relative to how processes do - or would - work in the NEM. These are discussed in Table 3.

Table 3 - Assumptions and simplifications

#	ASSUMPTION	IMPACT
1	Intra-regional constraints are simplified, with generating units modelled as being located across nine sub-regions, as used in the 2022 Integrated System Plan (ISP) capacity outlook model (see Figure 3). For each region, all load is allocated to the sub-region that contains that region's regional reference node (RRN).	Major transmission constraints are thought to be approximately captured with this approach, noting that the model used in the NEMDE is also an approximation to the real transmission system. The real NEMDE and PD processes use many more transmission constraints, and this would increase the number of variables and constraints – however, it is not clear that this would necessarily increase solution times as these variables and constraints should be continuous and linear, and therefore not computationally complex.
2	The last half hourly run in any PDS run after an OSM schedule has run is treated as defining the real-time (RT) outcomes, i.e., the actual NEMDE dispatch for settlement purposes.	NEMDE features such as five-minute dispatch or fast start unit activation are not included, however this would not affect the OSM itself, as it is upstream of NEMDE.
3	The existence of the OSM could impact participant contracting, changing the PD offers, which is ignored.	While this may affect results (e.g., which units are dispatched), this study is concerned with exploring how the OSM could be formulated and how it may work, and not with predicting actual future outcomes.
4	Realistic transmission loss factors are not included, although a small penalty on losses is used to avoid degenerate solutions, e.g., where otherwise there could be non-zero flows in both the forward and reverse directions of a single transmission link.	This will affect the relative economics of different generating units in the model, but this affect is a) thought to be small, and b) would not affect the primary objectives of studying how the OSM would work.

<table>
<tr><th>#</th><th>ASSUMPTION</th><th>IMPACT</th></tr>
<tr><td>5</td><td>Unless explicitly defined as part of a case study, we do not consider any exogenous events such as generator outages or forecast errors (for e.g., demand, wind, solar). In reality, such stochastic events may mean that OSM decisions are not optimal given actual conditions in real-time.</td><td>While the OSM should perform adequately given uncertainty, it is necessary to first explore the OSM formulation under a relatively controlled deterministic set of inputs. However, the scheduling parameter assessment in particular includes wind, solar and demand forecast errors.

Note that should conditions change such that the OSM solution is no longer secure, AEMO would still be able to issue directions.</td></tr>
<tr><td>6</td><td>Some simplifications are made to security constraints in the OSM.

• While for the Victorian and South Australian regions, acceptable security configurations have been published[3], these configurations were not available for other regions. Hence, for these other regions, a simple constraint requiring a minimum number of selected synchronous units to be running in each period is included.

• Minimum inertia constraints are included for some regions in some scenarios (but not included by default), taken from AEMO's 2022 Inertia Report. Some of these inertia values require certain levels of FFR to be available, however this FFR requirement was not included in the modelling.

Case studies focus on outcomes in South Australia and Victoria, these being the most interesting regions.</td><td>These configurations in SA and Victoria are representative of the current situation in the NEM, with other regions typically having enough synchronous generation online such that configurations for those regions have not yet been defined.

Therefore, the inclusion of minimum synchronous units is thought to a) be reasonably realistic and b) provide some value in 'testing' the OSM with different types of security constraints.</td></tr>
<tr><td>7</td><td>Certain gas and hydro units are able to contribute to inertia constraints via two modes:

• Synchronous generation (SG) mode, which is how a generating unit would normally run when generating energy into the NEM power system.

• Synchronous condenser (SC) mode, in which the machine is electrically connected to the power system, but without activation of the prime mover/turbine. This allows the machine to contribute to security constraints (provide inertia), at close-to-zero power.

When running in SC mode, it is assumed that the units provide less than 25% of the inertia provided in SG mode. A constraint was tested that requires a (configurable) set of units to undergo a delay in 'switching' between modes. This is described further in the formulation in Section 4.</td><td>Note that this assumption is only relevant to one scenario where it was assumed that SA is islanded and an inertia constraint is applied. This islanded case study was used because it allows examination of a separate service requirement for inertia, using publicly available inertia requirements.[4]

The purpose of this scenario is to test the constraint themselves, not to make assumptions about uptake and behaviour of synchronous condenser capability.

The assumption that less inertia is provided is thought realistic, and means there is a trade-off between committing a unit in SC mode at lower cost, or obtaining more energy but at a higher cost via SG mode.

Units in synchronous condenser mode were not able to contribute to a configuration.</td></tr>
</table>

[3] AEMO – Transfer Limit Advice – System Strength in SA and Victoria, September 2022

[4] The inertia requirement was set at 6,200 MW.sec when SA is islanded, as set out in AEMO – 2022 Inertia Report.

#	ASSUMPTION	IMPACT
8	The model does not explicitly calculate or optimise carbon emissions or any other environmental consideration.	It is possible that the additional emissions from committing a thermal OSM unit exceed the emissions saving from dispatching additional renewable (IBR) capacity. However, to the extent that environmental policy is reflected in offers, this would be implicitly included in the optimisation.

Finally, it is re-stated that in some scenarios, input data and assumptions are set for the purposes of testing the prototype and its possible behaviour, and not to make a judgement call as to whether those assumptions are likely to hold in the future. As the OSM is a relatively novel approach to security, it is important that the model be tested against a wide-range of plausible – though not necessarily likely – inputs.

Figure 3 - Sub-regions used in the model[5]

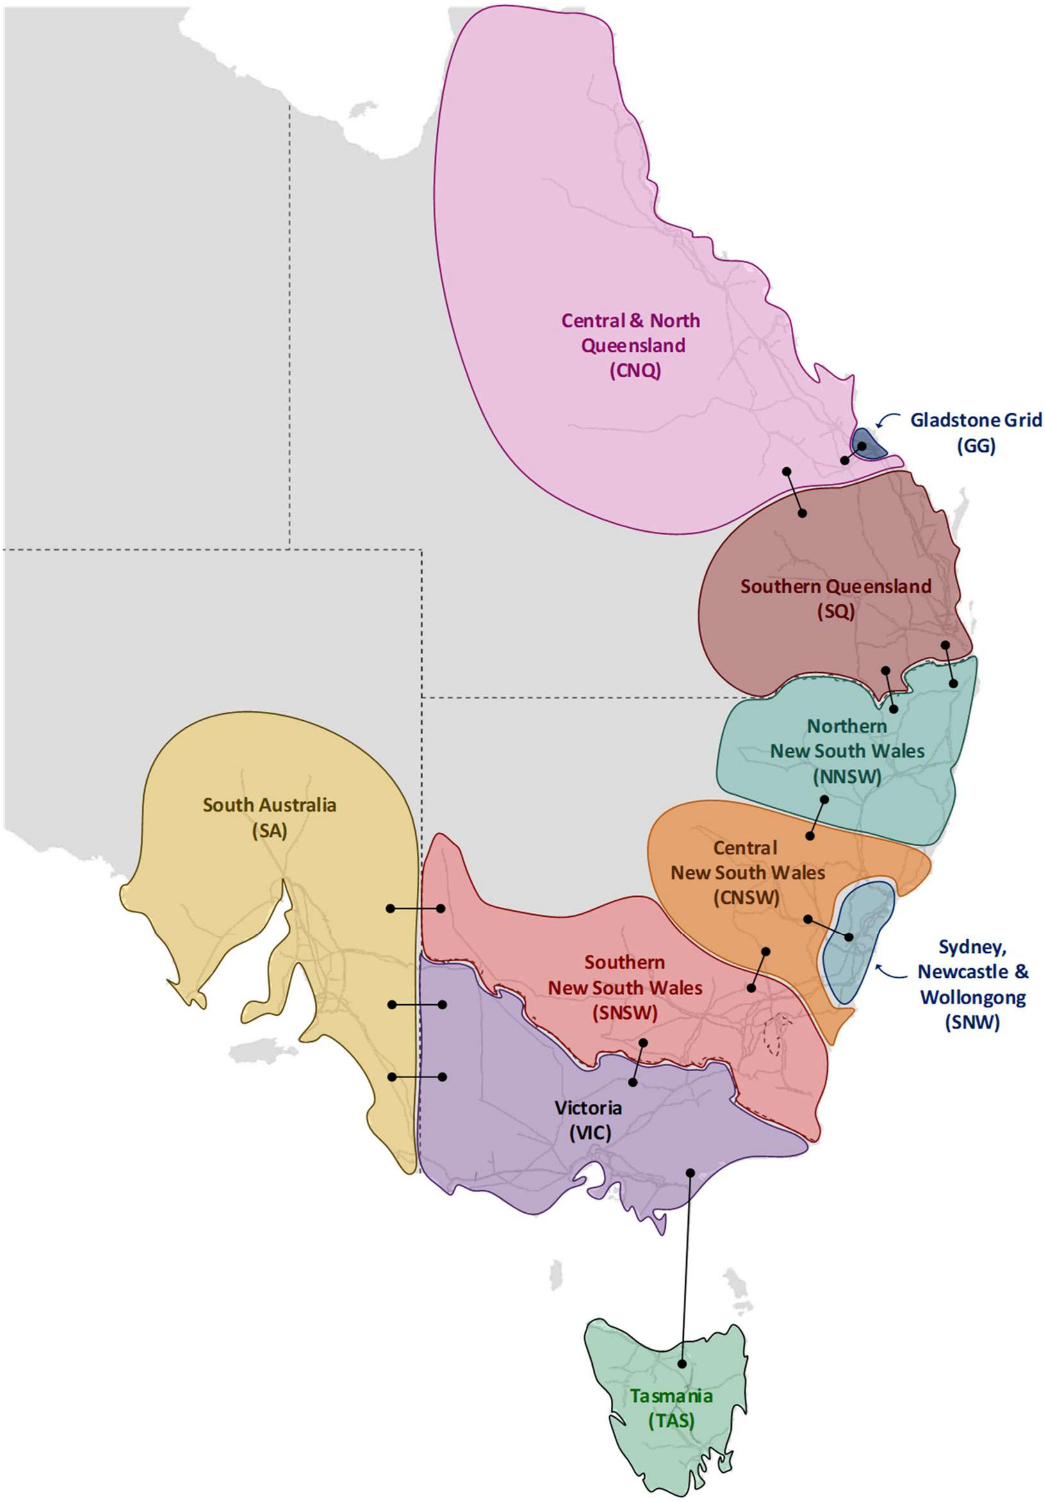

[5] Image from AEMO's 2022 ISP Forecasting Assumptions Update workbook.

3.2 Phases

While the focus of this work is to build and test a formulation for the OSM, it is also necessary to explore how that formulation interacts with centralised and decentralised market processes – for example running the PDS by AEMO, and updates to PDS bids by participants. Hence, the prototype used five separate models, referred to as phases as they are run sequentially. These are described as follows:

- Phase 1: Pre-OSM Pre-Dispatch. Phase 1 is a representation of the actual NEM pre-dispatch optimisation, in which units are scheduled based on their PD offers. The model is a continuous LP (e.g., no unit commitment or other integer/binary constraints are included, nor OSM security constraints). The output is a dispatch (and FCAS enablement) schedule, referred to as the PD schedule. Note that synchronous condensers are not scheduled in this phase, as they do not produce energy.

- Phase 2: Pre-OSM feasibility check. Phase 2 makes any required adjustments to the PD schedule from Phase 1 to ensure that it is feasible with respect to resource operating constraints, e.g., minimum on or off times and minimum generation. This is required as Phase 1 does not explicitly include these constraints, and are only met if generator offers are implicitly consistent with these constraints (which may not be the case). In the real NEM processes, this step would be undertaken by participants reviewing their units' PD schedule and revising offers as necessary.

- Phase 3: Operational Security Mechanism: Phase 3 commits/schedules OSM units in order to meet resource operating constraints and OSM security constraints, and otherwise maximise the gains from trade by scheduling additional security. This phase is the focus of this document and is discussed in detail below.

- Phase 4: Post-OSM offer adjustment: Phase 4 is a re-run of the Phase 2 PD with constraints or offer modifications to ensure that any OSM commitments are incorporated. As for phase 2, it applies UC constraints to ensure that schedules remain feasible for all units.

- Phase 5: Post-OSM PDS: As for Phase 1, but taking into account the OSM commitments as in Phase 4.

A key input to the phase 3 OSM simulation is the self-commitment flag for each unit in each period. For the purposes of the prototype, a self-commitment flag is determined for all units that can contribute to an OSM security constraint. When the flag is set to 1, the unit is assumed to be online and synchronised for that interval, and hence providing inertia and contributing to the selection of relevant security configurations. The self-commitment flag is set to 1 where a unit is scheduled to run at or above minimum generation in the phase 2 PD, and its PD revenue (using the phase 1 prices) exceeds its incurred fuel and start-up costs over the contiguous periods it is online.

However, interpretation of the self-commitment flag has not been finalised in the OSM design work, and may be used differently – e.g., it could be used to opt-out of consideration in the OSM, while carrying no implication that the unit does or does not intend to be online for the relevant intervals. The participant could then use the self-commitment flag where it is still considering whether to run its units (and wants to avoid an OSM commitment obligation). This would then mean that there may be units which do not run through a normal self-commitment, and which cannot be committed by the OSM.

Figure 4 - Modelling flow for OSM prototype toy model

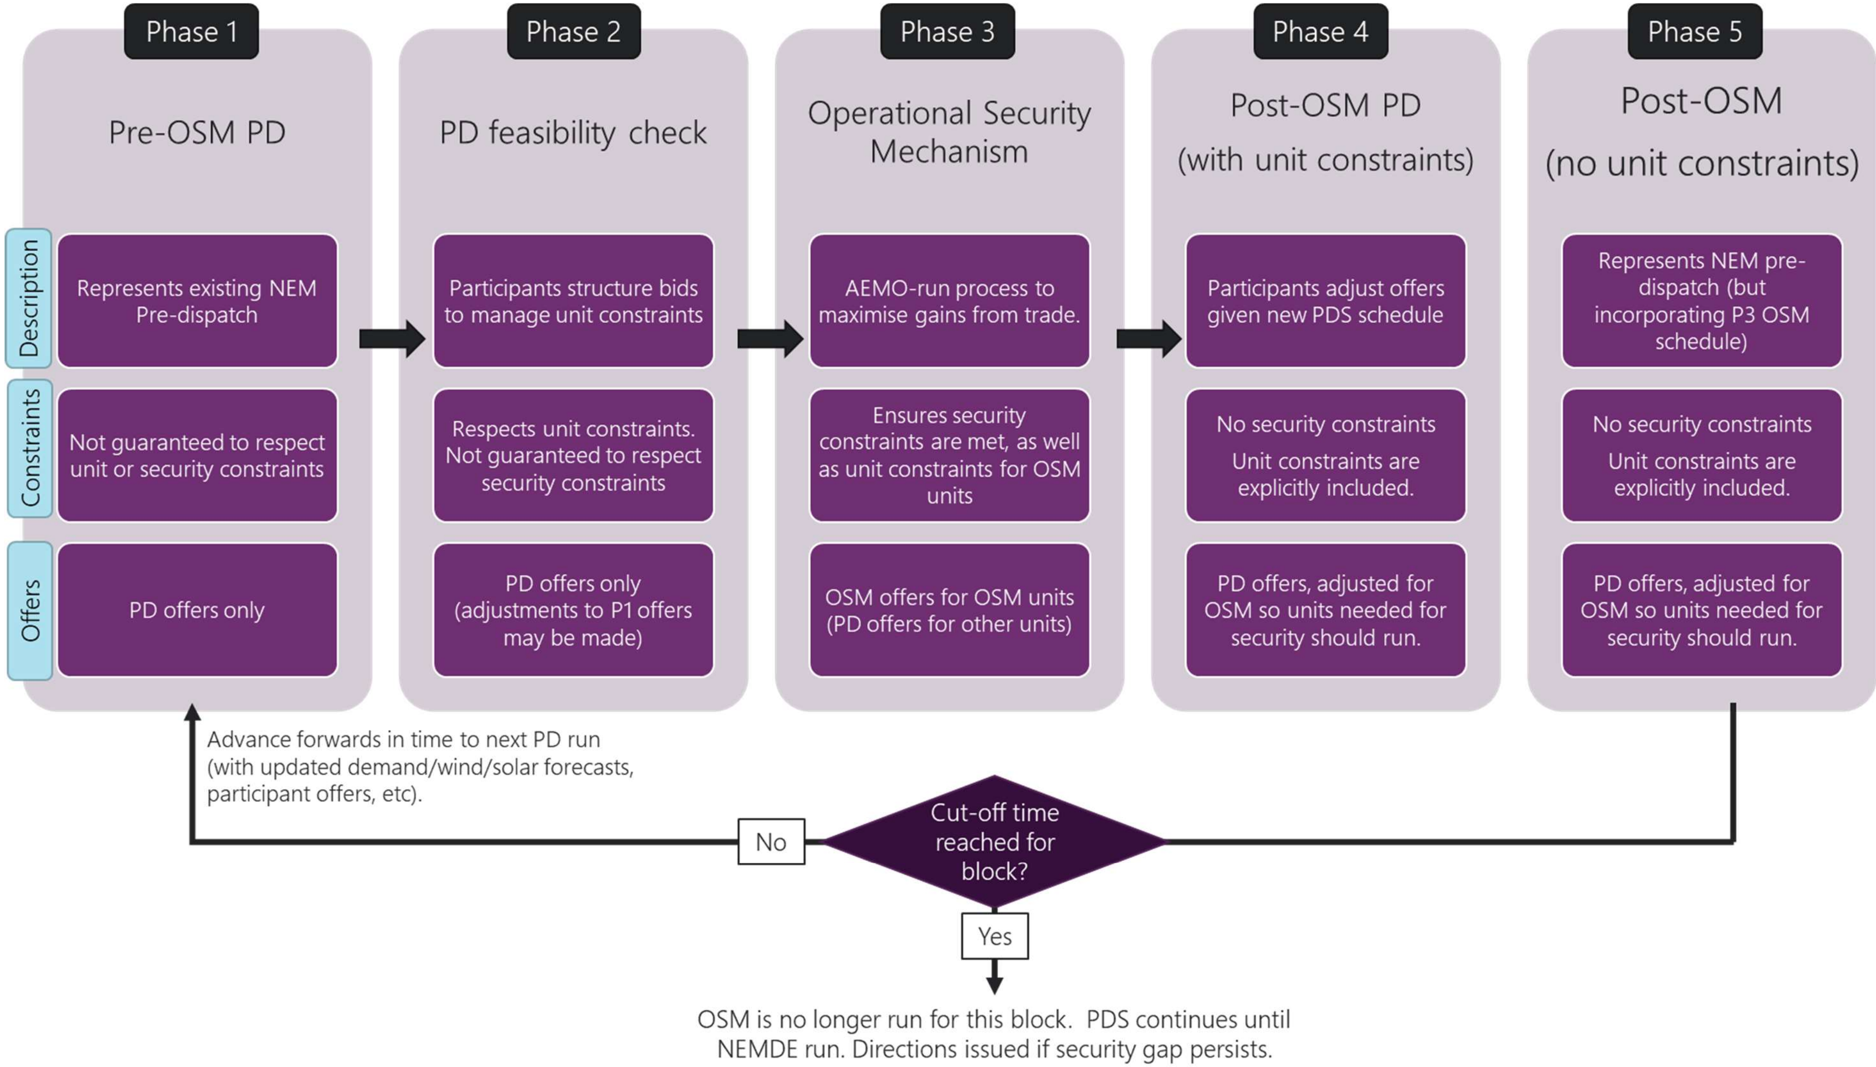

Table 4 presents the constraints included in each phase. A constraint is referred to as being reflected on the left-hand side (LHS) where the constraint limit it is part of a variable and can be changed by the optimiser. It referred to as being on the right-hand side (RHS) where the constraint limit is a fixed parameter and cannot be changed by the optimiser.

Table 4 - Constraints included in each phase

CONSTRAINT	PHASE 1	PHASE 2	PHASE 3	PHASE 4	PHASE 5
FCAS Requirements	Yes	Yes	Yes	Yes	Yes
Minimum On/Off Time	No	Yes	Yes	Yes	No
Minimum Generation	No[6]	Yes	Yes	Yes	No (phase 1 footnote applies)
OSM security constraints[7]	No	No	Yes	No	No
Maximum constraint on SA Wind and Solar[8]	Yes – 1,300 MW	Yes – 1,300 MW	Yes – LHS variable	Yes – RHS parameter updated from P3	Yes – RHS parameter updated from P3
Maximum constraint on VIC imports[9]	Yes – RHS parameter	Yes – RHS parameter	Yes – LHS variable	Yes – RHS parameter updated from P3	Yes – RHS parameter updated from P3

3.3 Processing Between Phases

3.3.1 Initial conditions

No specific initial conditions (e.g., initial units online, initial generation for each unit) are imposed, and the optimiser is free to for the phase 1 to be consistent with the unit's schedule from the run. Note also that storage volumes must be restored by the end of the day. This would not be the case in a production system, since the initial scheduled generation for each unit will be known and storage volumes will not be tracked.

3.3.2 Phase 1: PD price determination

The first phase provides an unconstrained baseline to the optimal scheduling of units and their dispatch with the only constraints being the energy balance constraints by subregion and the minimum FCAS constraints. The only output from this phase is the energy and FCAS prices for the five NEM regions. The energy price is determined from the shadow price of the energy balance constraint for the subregion representing each region's regional reference node. Prices for each FCAS in each region and interval are similarly determined using the FCAS requirement constraints.

[6] Although generation of FCAS may cause some minimum here
[7] Refers to a secure operating level of inertia, minimum number of synchronous units, and secure configurations.
[8] Applies in South Australia only.
[9] Applies in Victoria only. Note that while the limits are stated as being applied in phase 1 and 2, because the vast majority of configurations that currently apply do not have any restriction on imports, this constraint is set to a large number (effectively unconstrained). However, the functionality exists should a more restrictive number need to be tested.

3.3.3 Phase 2: identifying self-committed units

It is important to identify the units that self-committed, as these units cannot be committed by the OSM. In our modelling, these units are assumed to be online and contributing to security requirements. In reality, OSM participants would 'declare' their unit's self-commitment status by interval. The toy model tested two possible models for how participants set this flag:.

- In the first approach OSM units are taken to be self-committed if they have submitted a negatively priced offer with non-zero quantity in an interval, as this indicates their minimum stable generation is being offered to ensure it is scheduled.
- In the second approach, the revenue a unit would earn over its commitment period was compared with its incurred costs, and was taken to be self-committed if it made positive profit over this period.

The first approach was found to give better outcomes, and was used for all results presented in this report.

3.3.4 Phase 3: OSM commitments

Units that are committed in the OSM are identified, and these units must be scheduled in the following phases 4 and 5, either by submitting offers that are low enough to be scheduled, or through constraints (or both).

3.3.5 The 4th and 5th phases

For these phases, adjustments are made to the PD offers:

For OSM committed units, PD offers are revised so that a) their band 1 quantity is set to the unit's minimum generation, and b) the band 1 price is set to the price floor. Other bands are then adjusted to remove this minimum generation quantity, and to combine the last two price bands, which is needed because the addition of a band 1 price/quantity at the price floor (if such a band does not exist) would result in 11 bands.

Phase 4 then runs the PD with unit commitment constraints (as for phase 2), and phase 5 runs the PD without unit commitment constraints (as for phase 1).

3.3.6 Other features and limitations

The following limitations or assumptions apply to the current model:

- It is assumed that all units which have the capability to be accredited for the OSM have chosen to do so. In reality, some units may not be accredited (e.g., participant expects OSM revenues is too low), and the OSM could not commit these units. It is planned to include units which do not participate in the OSM (but have the capability to do so if they wished to) in the prototype in the future.
- The enablement cost OSM bid cost component is ignored in phase 3 if the unit's self-commitment schedule from phase 2 is such that the OSM does not cause it to undergo an additional start (e.g., if the OSM only extends its commitment). In this case, its enablement cost component would also be ignored in settlement. Another case is when the OSM causes an avoided-start, that is, the OSM adjoins two separate self-commitment events. In this case, the enablement cost is subtracted both from the objective function (a cost reduction) and in the participant's settlement equations.
- Design discussions have envisaged that PD offers would be updated in reaction to a changed OSM commitment (e.g., if one unit is committed, another may de-commit). This is not built into the model (due to the complexity of doing so), but in the future could be explored e.g., use a base case and focus on a few units that change bids to test logic around convergence.

3.4 Settlement

While settlement has not been implemented in the prototype as yet, it would use the following approach for OSM and non-OSM units.

OSM units are paid according to their cleared OSM offers:

- Variable price component ($/MWh or $/h):
  - For units that generate electricity when providing security services (e.g., hydro/thermal units), they are paid their variable bid price ($/MWh) against their running time and minimum generation when OSM committed. Any additional generation above this minimum is then paid the real-time spot energy price (described below), and receives no payment through the OSM.
  - For units that don't need to generate, e.g., synchronous condenser mode units and batteries, a $/hour bid component is used instead of a $/MWh.
- **Enablement price component ($):** OSM units are paid their start-up/enablement cost if their OSM commitment requires them to undergo an additional start-up process, relative to their self-commitment schedule. If the OSM only extends their commitment (e.g., unit is self-committed from 6 am through to 12 pm, and the OSM commits the unit from 12 pm to 6 pm) then no start-up cost is paid.[10]

OSM units may also earn additional spot market revenue for energy they provide above their minimum enablement. They may also earn FCAS revenue if enabled. These payments occur at the prices and dispatch schedules determined in the Phase 5 run (which is the final PD, but is also taken to be the real-time outcome).

As an example, consider a unit with a 20 MW minimum generation, and submitted start-up costs and variable costs of $1000, and $50/MWh respectively where each period is of one hour duration. Commitment is shown in Table 5.

[10] The draft determination also contemplates that revenue equal to the unit's enablement price may be deducted if the OSM commitment causes the unit to avoid a decommitment and subsequent re-commitment. For example, if a unit would be decommitting at 10 am, and re-committing at 2 pm, it would be incur a start-up cost in doing so. If the OSM then requires the unit to remain committed from 10 am to 2 pm, then it has avoided this start-up cost.

Table 5 - Settlement example data

PERIOD	COMMITMENT TYPE	DISPATCH (MW)	ENERGY PRICE ($/MWH)
1	Self-committed	50	$100
2	Self-committed	50	$100
3	OSM-committed	20	$5
4	OSM-committed	50	$60
5	OSM-committed	20	$5
6	Off	0	$5

In this case the unit will earn:

- **Periods 1 & 2 – Self-committed:** The unit is simply paid the energy price for two hours, that is [50 MW × $100/MWh × 2 hrs] = $10,000.
- **Periods 3 – 5 – OSM committed:** The unit is paid the variable cost on its minimum generation for three hours, that is [20 MW × $50/MWh × 3 hrs] = $3,000. The energy price is no relevant to this calculation. The unit is not paid its start-up cost, as it was already online in interval 2.
- **Period 4 – Generation above minimum enablement:** The unit is paid the energy price on its generation above its minimum enablement of 20 MW, that is [30 MW × $60/MWh × 1 hrs] = $1,800.

Total settlement across these periods is then $14,800, which exceeds the revenue of $13,200 if it were self-committed in all five periods.

While non-OSM units are not the focus of this study, settlement is also calculated for these units. As these units do not participate in the OSM, they are simply paid the prices from phase 5 against their energy and FCAS schedules in phase 5.

4 MODEL FORMULATION

4.1 Formulation overview

A brief overview of the formulation is provided to outline the key features, while abstracting away some of the mathematical details. Note that the term resource is used to refer to both generating units, and synchronous condenser.

4.1.1 Time intervals

There are two types of time intervals, which allows decoupling of commitment and dispatch decisions:

- Dispatch decisions are made over dispatch intervals *d*, which are half-hourly, as in the PD.
- Capacity commitment decisions are made over capacity intervals *c*.

The duration of the commitment interval can be varied. While it could be half-hourly (so that there is a one-to-one relationship with dispatch commitment intervals), it could also be longer. For example, if it were set to be two hours, there would be four dispatch decisions per commitment decision. A thermal unit that is committed for one commitment interval must then be dispatched above its minimum generation in each of the four dispatch intervals.

The effect of a longer commitment interval is to reduce the number of commitment decisions, which should result in improved computational decisions, at the cost of a potentially lower quality solution, because there is less freedom in making scheduling decisions.

4.1.2 Unit commitment modelling

Approaches to modelling unit commitment problems are well-documented in industry and academia, and the approach taken is relatively standard.

Binary variables for each resource and commitment interval are used to represent commitment status (1 if a unit is committed, zero otherwise). These can then be used to formulate energy output constraints, minimum on/off times, etc, e.g.:

- **Maximum generation:** dispatched power (MW) from a resource must be less than the product of its commitment variable and capacity (MW), which would be either its capacity or zero, if it is or is not committed respectively.
- **Minimum stable generation:** similarly, dispatched power must be greater than the product of the commitment variable and minimum stable generation (MW).

Similarly, additional binary variables are used to track whether units are starting-up or shutting-down, and whether they are operating in synchronous condenser mode.

4.1.3 Security constraints

Three types of security constraints are included, being i) the minimum number of synchronous units, ii) a requirement to be in a secure configuration, and iii) an inertia constraint. The inertia constraint is only included in certain scenarios.

- **Minimum synchronous units:** Certain units are classed as being synchronous (only hydro-thermal units). For each interval, the sum of the commitment binary variable of all such units in a region must meet or exceed the required number of units.
- **Secure configurations:** Inclusion of this constraint is more complex.

– For a region[11], there is a set of configurations, each of which has a binary variable representing whether it is selected in an interval. There is also a set of 'constraint station' groupings, for example, elements of this set are the Torrens Island power station and the Pelican Point power station. For each configuration, minimum numbers of units from each constraint station must be committed, e.g., configuration SA_70 requires to Torrens B units, and one Pelican Point unit. To implement this, the following constraints are used:

    - The sum over all regional configurations must equal one in each period, so one – and only one – configuration can be selected per period.

    - The configuration selection variable is used so that if the configuration is active, then for each constraint station, the number of units committed (sum over commitment variable for each applicable unit) must meet or exceed the number in that configuration. This is done by multiplying the minimum units value by the constraint selection variable (which is zero if a configuration is not selected, effectively removing the minimum unit requirement).

- **Inertia:** The inertia from a unit is the product of its inertia (a fixed value) and its commitment variable. The sum over all units' inertia (in a region) must meet or exceed the regional inertia requirement, when applicable.

Note that the minimum number of units for an active configuration is formulated as a greater-than-or-equal-to constraint, not an equal-to constraint. Therefore, SA_70 which requires two Torrens B units and one Pelican Point unit, could be active if there are four Torrens B units and one Pelican Point unit online.

4.1.4 Limits on IBR due to a selected configuration

Beyond the need to be in a secure configuration for minimum security requirements, there may be benefit in incurring additional costs to select a configuration that relieves limits on power system operation, e.g., allowable IBR. For example, configuration SA_30 allows up to 1,900 MW of SA IBR to be dispatched, while SA_31 allows 2,000 MW. Both are satisfactory from a minimum-security perspective, but the latter may be preferable (even if more expensive) if the additional 100 MW of IBR significantly reduces overall costs.

To implement these limits, a constraint is applied on total SA IBR, limiting it to the sum of the product of the configuration selection variable and allowable IBR in each configuration. This also shows why only one configuration can be selected, even though in reality enough units may be committed for multiple configurations – otherwise the allowable IBR would be summed.

4.2 Avoiding energy-only OSM commitments

4.2.1 Description of the issue

The OSM formulation makes decisions to commit OSM units to maximise the gains from trade, subject to also meeting minimum security requirements. In a typical scenario, committing an OSM would tend to increase dispatch costs, but the gains from trade are *indirectly* improved through such a commitment because it relaxes constraints on (low-cost) generation from IBR. In the future, it is possible that the OSM evolves such that other types of constraints might be relaxed by an OSM commitment.[12]

However, there is another scenario in which a committed OSM unit *directly* improves the gains from trade, simply because its OSM offers mean it would supply energy at lower cost than the marginal unit. As an example, consider an OSM unit with OSM offers of a $4,000 enablement cost, and $35/MWh variable cost, and a minimum generation of 100 MW. If the unit were run for four hours, this gives it an average cost of energy of $45/MWh. If the marginal unit has a cost above this – say $50/MWh – the gains from trade are improved directly from committing this OSM unit, regardless of whether it serves any security purpose. This is referred to as an energy-only commitment,

[11] Currently SA and Victoria only, but in principle would apply to any region.
[12] As an example, a small number of secure configurations that apply in Victoria restrict imports from other regions.

distinct from a security commitment either to meet minimum security requirements, or obtain additional security that improves the gains from trade (or both).

Importantly, based on the OSM draft determination, the OSM should not make energy-only commitments, i.e., it should commit units for the benefit of the security provided, not solely for their energy. The remainder of this section presents a proposed solution and worked examples to give effect to the desired outcomes.

As an aside, it appears that energy-only OSM commitments should not typically occur, due to bidding incentives. If the energy price is high enough that a unit is economic to run directly, it would be expected to be dispatched through the PD itself, because if it is committed through the OSM, it is likely to have some excess revenue deducted in settlement. That said, there may be certain advantages in being committed through the OSM, namely that the unit could receive a revenue guarantee for its OSM costs that it would not receive when it directly faces the uncertain energy price. Ultimately, it is not possible to predict how participants would offer, and it is important that the formulation take these issues into account.

4.2.2 Proposed solution

A solution is proposed in which two runs of the OSM are used to arrive at the final schedule of OSM commitments. Note that this has not been implemented in the toy model, and therefore results shown in this report do not use the approach described below.

The two-run approach is:

- **OSM Run 1:** The OSM is run as it has been described above (and in Appendix A), using the OSM offers submitted as-is by the relevant participants. As described above, it is possible that some OSM commitments would be made as an energy-only commitment, i.e., with no net improvement to security.

- **OSM Run 2:** The OSM is then re-run to remove any energy-only commitments, while still meeting security constraints and committing the necessary units to be able to dispatch the efficient level of IBR generation. To do this:

  – All OSM offers are increased (e.g., multiply prices by a large number) so they would never be economic to run for their energy alone.

  – A constraint is added such that the total dispatched IBR (in SA) must be at least as much as that dispatched in Run 1.

  – Any OSM units that were not dispatched in Run 1 are removed from the optimisation (but only for the intervals they were not needed). This is because these units should not be scheduled in these intervals. This step is optional, because if the offer prices of all units are increased logically, these units would not be scheduled anyway.

Essentially, this second step a) makes the OSM units so expensive that they will never be committed unless required for a constraint b) turns the efficient level of IBR found in step 1 into a constraint that must be met by the optimiser. Note that if configurations can give rise to other changes in RHS limits (for example, some Victorian configurations include limits on imports into Victoria), these should also be included as constraints in Run 2.

4.2.3 Worked example

Setup

We consider a scenario in which there is a single region with some IBR and five units participating as in Table 6. Further:

- There is a security gap for three hours;

- The marginal supplier has a cost of $80/MWh;

- There is 1400 MW of IBR available, at a marginal cost of $5/MWh.

The table therefore shows the total, and average cost of energy for 3 hours for each unit, with A being the cheapest. Note that all units except E are cheaper than the marginal energy price.

Table 6 – Unit data for energy-only commitment worked example

UNIT	ENABLEMENT COST ($)	VARIABLE COST ($/MWH)	MINIMUM GENERATION (MW)	COMMITMENT COST ($)	AV. COST OF ENERGY ($/MWH)
A	$3,000	$35	100	$13,500	$45
B	$3,000	$45	100	$16,500	$55
C	$0	$60	100	$18,000	$60
D	$0	$70	100	$21,000	$70
E	$0	$150	100	$45,000	$150

Table 7 also shows the three configurations, their IBR and total cost over three hours, with C1 being the cheapest, but more expensive configurations allowing more IBR.

Table 7 - Worked example configurations

#	CONFIGURATION	IBR LIMIT (MW)	MINIMUM COST (3 HOURS)
C1	A, B	1000	13,500 + 16,500 = $30,000
C2	A, B, C	1500	13,500 + 16,500 + 18,000 = $48,000
C3	A, B, D	1800	13,500 + 16,500 + 21,000 = $51,000
C4	A, B, E	2000	$13,500 + $16,500 + $45,000 = $75,000

Run 1

Units would be committed if they meet at least one of three criteria (see also Table 8):

- **Needed for security:** Both A and B are needed required. Other units are not.
- **Gains from trade improvement via OSM energy:** Where the OSM offers are such that the average cost of energy is less than the marginal supplier, it would directly improve the gains from trade. Marginal cost is $80/MWh, so for each unit, a minimum generation of 100 MW for three hours saves 3hrs × 100 MW × $80/MWh = $24,000 in the objective function. Therefore, units A, B, C and D would meet this criterion.
- **Gains from trade improvement via IBR:** This occurs when the cost of an OSM commitment is less than the value of the additional IBR. In this case C allows an extra 500 MW - though only 400 MW is available. Typically, the cost of committing C would be compared against the difference between the marginal energy price and the IBR offer cost. However, this is unnecessary as C as economic as in the previous bullet. Similarly, D would also allow for extra IBR, but only if C is not committed.

Table 8 - Consideration of the need for each unit

UNIT	NEEDED FOR SECURITY?	DIRECTLY IMPROVES GFT	EXTRA IBR IMPROVES GFT?	COMMITTED?
A	Yes	Yes	NA	Yes
B	Yes	Yes	NA	Yes
C	Possibly	Yes	Yes (+ 400 MW)	Yes
D	Possibly (w/o C)	Yes	Possibly (+ 400 MW w/o C)	Yes
E	No	No	No	No

In sum, we have A, B, C and D online in the Run 1 solution, and the optimiser could choose either C1 or C2 because both allow all IBR to be dispatched.

Run 2

Now in Run 2, the OSM offers are scaled by – for example – a factor of 100. Now the cost of committing C becomes $1,350,000, and no unit will directly improve the gains from trade. The constraint requiring at least one configuration to be chosen still applies.

An extra constraint requires at least 1,400 MW of IBR to be dispatched, this being the efficient level of IBR. Without this, Run 2 would schedule only A and B, selecting configuration C1. However, this allows only 1000 MW of IBR, and so either C or D must be committed. C is cheaper, so C2 is selected.

Variations and special situations

The following points consider how the proposed approach would behave in some variations and special situations.

- In the worked example, if there was 1,850 of IBR available, then this would be scheduled up to 1,800 MW in Run 1.[13] Then in Run 2, 1,800 MW of wind is required, and unit D must be scheduled to select C2. It is interesting to note that it is now acceptable to commit D and obtain it's cheaper energy, simply because there is more IBR available. In other words, unit D is technically still an energy-commitment, but it is not an energy-only commitment, because it is also providing enhanced security that allows more IBR.
- Consider a scenario where only one of two energy-commitments (X and Y) must be chosen in Run 2 for a three-hour security gap. If X and Y have minimum run times of 4 hours and 3 hours respectively, and X is cheaper for four hours than Y is for three hours, then X would still be chosen. In other words, it may be an energy-only commitment for hour 4 in isolation, but is not an energy-only commitment when viewed over the whole four hours. Note that if it were technically feasible to run unit X for the first three hours only (no minimum run time), then Run 2 would give effect to this.
- As an extension to the previous point, consider a scenario where a unit could be committed for three hours to dispatch additional IBR, however, its enablement cost is slightly too expensive to make this economically efficient. But, if in the 4th hour, there is no IBR available, and there is a very high price spike, it may then be worthwhile to start the unit in Run 1. In this case, the unit is not an energy-only commitment, but nor is it a security-only benefit – the costs of committing the unit come partially from the security value, and partially from the energy value, and neither alone is sufficient. In this case, the proposed solution would commit the unit.

[13] Unit E would not be committed to dispatch the extra 50 because it costs $45,000 to run for three hours, and the additional 100 MW wind costs $5×150 MWh = $1,500. Meanwhile, 450 MWh is displaced (that costs $80/MWh), saving only $36,000.

- Another scenario occurs where a unit X is needed from say 8am to 12pm, and again from 5pm to 9pm. If it has a large start-up cost, it may be cheaper to keep it online from 12pm to 5pm. Like the second bullet, it may be an energy-only commitment for certain hours, but not over the whole commitment period.

Finally, it is worth re-stating that that the settlement outcomes would deduct additional revenue on the unit's minimum generation above total OSM costs, so that there is an incentive not to get scheduled via the OSM if the energy price is expected to be high enough to recover total costs.

5 DATA

5.1 Introduction

This section describes the data used in the OSM, including the criteria for selecting particular days as a basis for the case studies.

5.2 Analysis of suitable days

Data was initially obtained for the period of October 2021 through to December 2022. Demand wind and solar profiles across this period were analysed to identify suitable days which could be used as a basis for developing scenarios. Two types of days were identified:

- Low net-demand days (i.e., demand less available wind and solar), are expected to conditions in which – based on spot market energy prices alone - it is not economic for many synchronous resources to run, meaning security requirements may not be met.
- High IBR days with high demand are expected to have conditions in which there is value in relieving constraints on IBR to displace the higher cost resources that would otherwise be needed, in other words, to improve the gains from trade.

5.2.1 Low net-demand days

For each region, days were ranked based on a) lowest half-hourly net demand and b) lowest total net-demand (i.e., the sum of net demand in the 48 half-hourly periods). Note that operating days were used, i.e., starting and ending at 4 am. From these rankings, days which have relatively low demand in all or at least most regions were identified, concentrating particularly on South Australia and Victoria.

From this analysis, it was found that:

- 27th Dec 2021 was a low demand day for South Australia, Victoria and New South Wales
- 14th November 2021 was a low demand day for all regions.

These dates are to be used as the initial basis for developing case studies, though others may also be used if necessary for the OSM testing.

5.2.2 High-IBR days

Analysis was focused on IBR in SA, and demand in SA and neighbouring states. First, days with SA IBR exceeding 2,300 MW were selected. These were then examined for days in which SA demand exceeded 2200 MW, or there was high demand in Victoria or other regions:

- 22nd Aug 2022 was found to meet these criteria, with very high demand in NSW, and moderately high demand in Victoria and Queensland.
- 25th July 2022 was also found to have very high IBR, with reasonably high SA demand.

However, the July date was not preferred because the NEM market suspension state had only just been removed at this time, and this may have resulted in some non-typical outcomes in PD offers and other data.

5.3 Data for scenarios and design questions

Table 9 shows the base case simulations that are used to explore basic functionality of the formulation and OSM design.

Table 9 - Base case scenarios

#	SCENARIO OR QUESTION	DATE USED	MODIFICATIONS
BASE CASES			
1A	OSM units are brought on to improve gains of trade in South Australia by releasing more IBR.	22nd August 2022	PD offer prices from one Torrens Island B unit are increased so that it would not be scheduled in phase 1 and phase 2, meaning the system is not in a secure configuration without the OSM schedule. A variation increases available IBR, so that there is value in selecting a more expensive configuration.
1B	OSM units are brought on to satisfy a secure configuration and an inertia constraint in South Australia.	27th December 2022	SA assumed to be islanded, so that a 6,200 MW.sec inertia constraint applies in that region.
1C	OSM units brought on to meet a minimum sync unit constraint in Tasmania.	27th December 2022	The units which can contribute to this constraint were assumed to be:[14] • John Butters (hydro) • Gordon (hydro) • Catagunya (hydro) • Poatina (hydro) • Tamar Valley (gas combined cycle)

Table 10 then shows the approach to answering various research or design questions with the model or with worked examples.

[14] Note that these units were chosen based on their capacity, and it is understood that in reality these are not the only units that can place the Tasmanian system in a secure state. However, this is thought suitable for the purposes of testing the formulation.

Table 10 - Overview of research questions and approach

#	QUESTION	DATE OR BASE CASE USED	MODIFICATIONS
2	Should the OSM mechanism use bidding at the price floor in PD as an approach to ensure resources cleared in the OSM are scheduled in dispatch?	Base case with adjusted offers to be at the price floor	Offers modified to examine a case where the PDS price is at the price floor, to demonstrate the need for some mechanism to ensure that the security critical resources (whether self-committed or selected in the OSM) are committed.
3	What are the mechanics of the OSM relieving IBR constraints for net market benefit?	As for 1A	This focusses on South Australia, and contrasts two cases: • the commitment is to meet a secure configuration (perhaps the allowable IBR constraint is not binding) • more expensive units are chosen to allow for more IBR to be dispatched.
4	What are the mechanics of the OSM unbundling system services from configurations?	As for case 1B	As for case 1B. This is also shows the OSM using different configurations to schedule units when SA is (assumed to be) islanded.
5	How does the OSM treat units that can operate in multiple modes (e.g., fast-start units that can operate in synchronous condenser mode)? Some units can operate in synchronous condenser mode, but produce a lower value of inertia compared to operating in generation mode.	As for case 1B	
6	How should the economic benefit of wind/solar units that are disorderly bidding due to a constraint (e.g., bidding at floor price instead of approximate short-run cost) be treated?	Worked example developed demonstrate the precise outcomes.	

5.4 General Data

Table 11 describes types and sources for data used in the cases studies.

Table 11 - Data used in the case studies

DATA TYPE	SOUCRE	COMMENT
TRACES		

DATA TYPE	SOUCRE	COMMENT
Demand	Half-hourly operational demand by region, obtained from NEMWEB.[15]	
Solar forecast	Unconstrained Intermittent Generation Forecast (UIGF) obtained from NEMWEB.	The last submitted POE50 (50% probability of exceedance forecast) was used.
Wind forecast	UIGF obtained from NEMWEB.	The last submitted POE50 (50% probability of exceedance forecast) was used.
SECURITY REQUIREMENTS		
Secure Operating Level of Inertia	AEMO - 2022 Inertia Report	Applied on in Tasmania and South Australia (if the latter was assumed islanded, as for base case 1B).
Minimum Synchronous Units	AEMO - 2022 ISP Forecasting Assumptions Workbook	
Secure configurations (Victoria and South Australia)	AEMO - Transfer Limit Advice – System Strength in SA and Victoria	
Regulation FCAS Requirements	AEMO - Regulation FCAS Changes	180 MW raise and 170 MW lower
GENERATOR DATA		
PDS Bids	Obtained from NEMWEB	For each day, the last submitted energy offers were used. FCAS offers were obtained from 17th Oct 2022 only, and this was used for all simulated days (as FCAS is not the focus of the OSM).
OSM Bids	Start-up costs obtained from Acil Allen – Fuel and Technology Cost Review Running costs (at minimum generation) obtained from AEMO - 2022 ISP Forecasting Assumptions Workbook	
Unit commitment and other technical data	2022 ISP Forecasting Assumptions Workbook and NEM Registration and Exemption List	Includes capacity, ramp rates, minimum on and off times, and minimum generation.[16]
Inertia	Determined from technology based inertial constants, and shown in Table 12.	A selection of identified candidate hydro and gas units were considered able to operate in synchronous condenser mode. When operating in synchronous condenser mode, a significantly reduced inertial constant was applied to these units (approximately one quarter of the normal inertial constant).

[15] http://nemweb.com.au/

[16] Minimum on and off times were supplemented with data from the ACIL Allen Fuel and Technology Cost review (to provide more variation in times, as the ISP used a single for all relevant units)

DATA TYPE	SOUCRE	COMMENT
		In some cases, batteries are assumed to be capable of providing synthetic inertia. It is not clear what sort of operational constraints would apply for a battery to do so – for the purposes of this model, it is simply assumed that to do so the battery must reserve X% of its capacity as both headroom and footroom, with X in the range of 1% - 10%. In this state, it is then modelled as having an inertial constant of approximately 11 MW.sec.
FCAS Trapezium	Trapezium were assumed to have 45° upper and lower angles.	
Batteries and Pumped Hydro technical data	AEMO - 2022 ISP Forecasting Assumptions Workbook	Covers storage volumes and charge, discharge and pumping efficiencies. Inertia associated with a battery operating in "virtual machine mode" was obtained from Neoen – Virtual Machine Mode Test Report.[17]
TRANSMISSION SYSTEM DATA		
Regional Topology	AEMO - 2022 ISP Forecasting Assumptions Workbook	The five NEM regions were modelled, with nine sub-regions as in the 2022 ISP database.[18]
Flow limits	AEMO - 2022 ISP Forecasting Assumptions Workbook	

Table 12 shows the value of inertia constants for each technology and mode.

Table 12 - Inertial constants

TECHNOLOGY	INERTIA CONSTANT (MW.SEC)
Black Coal	3.44
Brown Coal	3.66
CCGT	6.90
OCGT	7.67
Liquid Fuel	4.00
Gas-powered steam turbine	4.70
Hydro	3.65
Pumped Hydro	5.44
Gas - reciprocating	1.58

[17] Neoen – Hornsdale Power Reserve Expansion Virtual Machine Mode Test Summary Report, 9th March 2022

[18] Victoria, South Australia and Tasmania use only a single sub-region each, while for Queensland and NSW, sub-regions are:

- Queensland: Central and North Queensland, South Queensland
- New South Wales: Northern NSW, Central NSW, Sydney Newcastle and Wollongong, Southern New South Wales

The South Australian synchronous condensers are assumed to provide 1,100 MW.sec of inertia. Their commitment status was exogenous to the model, and in all presented scenarios two were assumed to be committed.[19] For units operating in synchronous condenser mode, they were assumed to contribute an inertia quantity that is 25% of their normal mode value. Note also that when in synchronous condenser mode, it is assumed that units don't contribute to a configuration.

Wind, solar and battery technologies do not provide any inertia, except in specific testing of one battery in South Australia being able to operate in a mode that does so. In this case, it was assumed that the battery must reserve a proportion of its capacity in both the raise and lower directions, in which case it provides a fixed quantity of inertia. Certain values were assumed for testing purposes - this was to test that the constraints themselves work rather than to assess the value of battery inertia to the system, and as such are not reported on. There was no requirement to have a certain amount of stored energy (or ability to store additional energy) in order to provide inertia.

[19] AEMO – 2022 Inertia Report, Dec 2022.

6 CASE STUDY RESULTS

This section first presents results from the three base case runs, focusing on the Case 1A in which the gains from trade are improved by an OSM commitment. Next, a series of policy questions are discussed using observations from these and other cases.

6.1 Base cases

6.1.1 Case 1A: gains from trade

This case study demonstrates a case in which a gas unit is committed, improving the gains from trade by relieving constrained IBR. It uses the high-IBR day (2nd August 2022) as a basis. Note that it is assumed that two SA synchronous condensers are running, and the system is not islanded.

Figure 5 and Figure 6 show the energy schedules from the initial PD (phase 2) and OSM (phase 3) for South Australia. A 1,300 MW limit is applied in the PD, being the lowest allowable IBR in any configuration, and this is binding as SA wind availability is around 1,700 MW for this day. Day time SA demand is in the range of 2,000 – 2,500 MW, and South Australia is importing energy from Victoria for most of the day (note that if total supply exceeds shown SA demand, this indicates SA is exporting to Victoria).

For most of the day, there are no synchronous units online in South Australia in the PD schedule (and those that are were not considered to be self-committed under the requirement to submit a negative offer). Therefore, SA is never in a secure configuration.

Hence, when the OSM is run, additional units must be committed. A Pelican Point unit and a Quarantine unit (both gas units) are committed, which – along with the two exogenous SA synchronous condensers – active configuration SA_33. Doing so allows up to 1,900 MW of IBR generation, which is shown as increased wind in Figure 6, and which results in a significant decrease in SA imports. This allows marginal generators in other regions to be backed off, as discussed below. Note that the two gas units running at their minimum generation are also visible.

Figure 5 - Case 1A: Initial PD energy schedule

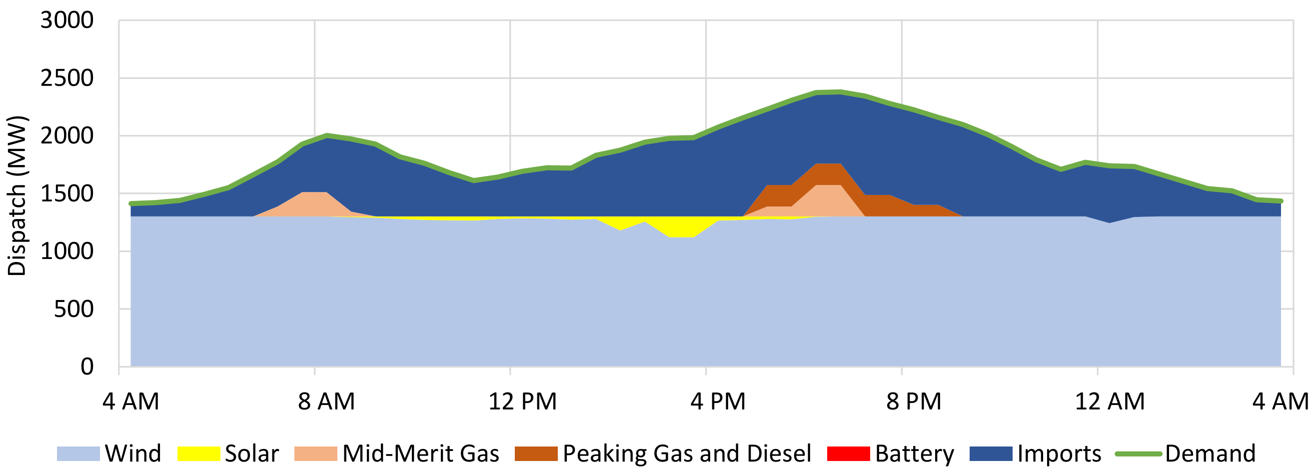

Figure 6 - Case 1A: OSM energy schedule

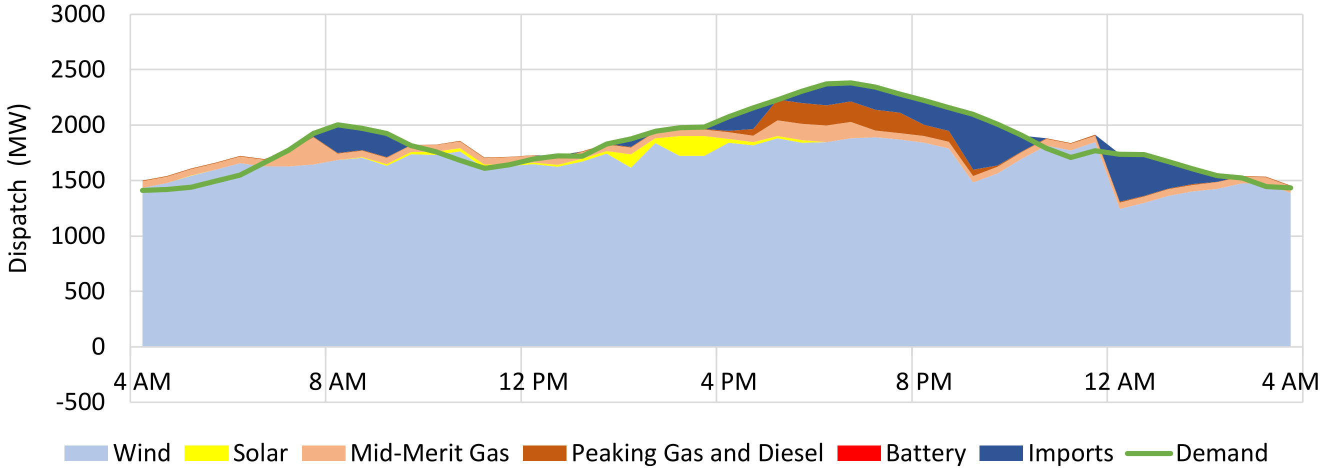

Table 13 shows the change in daily energy by technology and region, as a percentage of total NEM-wide daily demand. As mentioned, SA wind is dispatched upwards as a result of the relieved IBR constraint. This has the predominant effect of reducing coal output in Queensland, New South Wales and Victoria.

In addition, there is an increase in energy from Tasmanian hydro, of 0.6%. This occurs in order to meet the applied minimum large synchronous units constraint (see section 6.1.3), with two units - John Butters and Poatina – being run at minimum generation of 155 MW in addition to a unit that was running in phase 2. There is also increased OCGT production in Queensland. This is from a small station which erroneously had a low variable cost as fuel cost data was not available within the ISP dataset. It is thought that if a more realistic fuel cost was used for this station, it would not have been committed in phase 3 – however, this is a good example of how if OSM bids are not aligned with (lower than) the unit's PD offers, the second OSM run (as in Section 4.2) is required to avoid an energy-only commitment.

Table 13 - Case 1A: Change in daily energy by technology and region (percentage of daily demand)

TECHNOLOGY	NSW	QLD	SA	TAS	VIC	TOTAL
WIND	-	-	1.7%	-	-0.3%	**1.3%**
SOLAR	-	-	-	-	-	**-**
COAL	-0.9%	-0.4%	-	-	-1.1%	**-2.5%**
MID-MERIT GAS	-	-0.2%	0.3%	-	-	**0.1%**
PEAKING GAS AND DIESEL	-	0.4%	-	-	-	**0.4%**
BATTERY	-	-	-	-	-	**-**
HYDRO	-	-	-	0.6%	-0.2%	**0.5%**
TOTAL	**-0.9%**	**-0.2%**	**2.0%**	**0.6%**	**-1.6%**	

As demand is fixed (both in price and quantity), an improvement in the gains from trade occurs only through a reduction in supply costs, shown in Table 14. Costs are separated between those incurred via dispatch of PD offers and OSM offers, with the latter being zero in Phase 2 as the PD cannot make OSM commitments. $484,000 of OSM costs are incurred, but this is more than offset by $4.4M decrease in PD costs, so that the OSM schedule has higher gains from trade (as well as being secure).

Table 14 - Case 1A: Phase 2 and 3 dispatch costs ($000s)

QUANTITY	PHASE 2 (PD)	PHASE 3 (OSM)	DIFFERENCE
DEMAND VALUE	$566,206	$566,206	$0
PD OFFER COSTS	-$419,362	-$423,769	-$4,407
OSM OFFER COSTS	$0	$484	$484
TOTAL COSTS	-$419,362	-$423,285	-$3,923
GAINS FROM TRADE	**$985,568**	**$989,491**	**$3,923**

Relieving the IBR constraint has the effect of increasing available supply while holding demand constant, which can be expected to result in decreased prices. Indeed, Figure 7 shows the change in RRP for South Australia and Queensland in Phase 2 and Phase 3, where the morning and evening price spikes disappear in the latter. Note that the price in NSW is generally close to the price in Queensland, while the price in the Victorian and Tasmanian regions is generally aligned with that of South Australia, and hence are not shown.

Figure 7 - Case 1A: Dispatch prices in South Australia and Queensland.

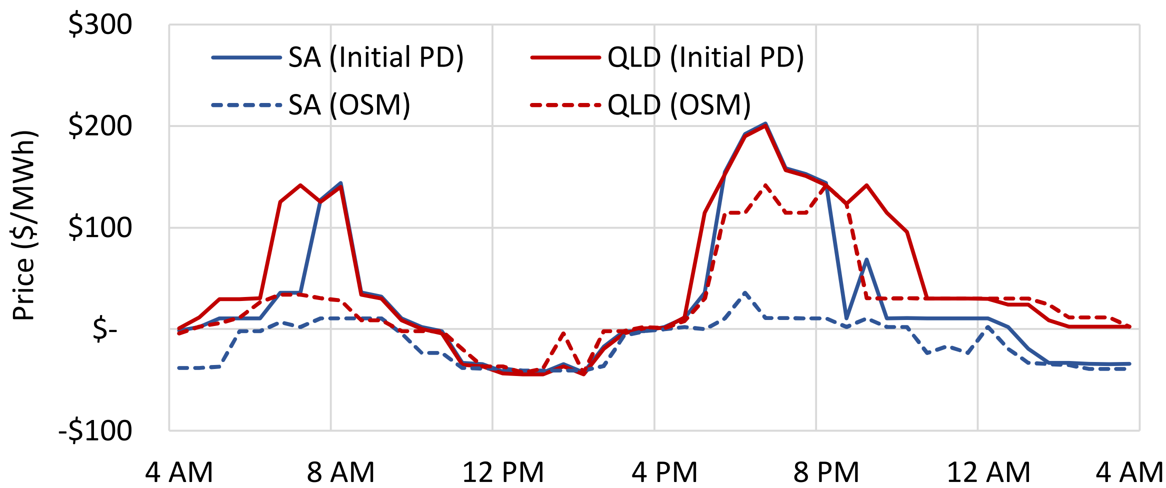

In particular, the morning and late afternoon price spikes are diminished in all regions. Notably the price falls from close to $200/MWh to close to $0/MWh in the late afternoon for the South Australian region. This price decrease is as a result of the increased efficiency from relieving power system constraints, but does also produce a wealth transfer from producers to consumers. Such dramatic changes in prices could be expected to result in changes to participant self-commitment schedules, with some units perhaps opting to decommit.

This scenario is somewhat contrived because with two SA synchronous condensers online, there is no secure configuration that allows less than 1,900 MW, whereas the real data had less than 1,800 MW of IBR available. Therefore, meeting the minimum-security requirements is sufficient to dispatch all IBR in this case.

To truly test whether the formulation would commit additional OSM units (with more cost) to dispatch more IBR, available wind was increased so that total IBR was greater than 2,000 MW (which is the maximum amount dispatchable) with two synchronous condensers online). Now the OSM chooses configuration SA_38, which requires three units (Torrens Island B, Mintaro, Quarantine gas units), the two synchronous condensers, and allows for 2,000 MW of IBR. Table 15 compares the OSM costs and gains from trade with the original and increased IBR available. OSM costs are increased, but the reduction in PD costs more than offsets this.

Table 15 - Case 1A gains from trade comparison ($000s)

QUANTITY	PHASE 3 ORIGINAL IBR	PHASE 3 IBR INCREASED > 2,000 MW
DEMAND VALUE	$566,206	$566,206
PD OFFER COSTS	-$423,769	-$429,539
OSM OFFER COSTS	$484	$526
TOTAL COSTS	-$423,285	-$429,013
GAINS FROM TRADE	$989,491	$995,219

6.1.2 Case 1B: inertia constraint

This case examines an inertia constraint that is applied to South Australia when it is in an islanded state (and hence has no imports or exports). The 27th December 2021 day is used, and it is assumed that two of the SA synchronous condensers are operating.

Figure 8 shows the PD energy schedule for South Australia. The rooftop solar carve out is visible in the middle of the day (distributed generation is not included in the model). In the middle of the day, wind generation is curtailed (for economic reasons) while the gas units (two Torrens B units and an Osborne unit) are operated at constant load. This indicates that the gas PD offers are lower in the merit order than the PD wind offers, which is required in order to respect their unit commitment constraints.

Figure 8 - Case 1B: SA dispatch schedule

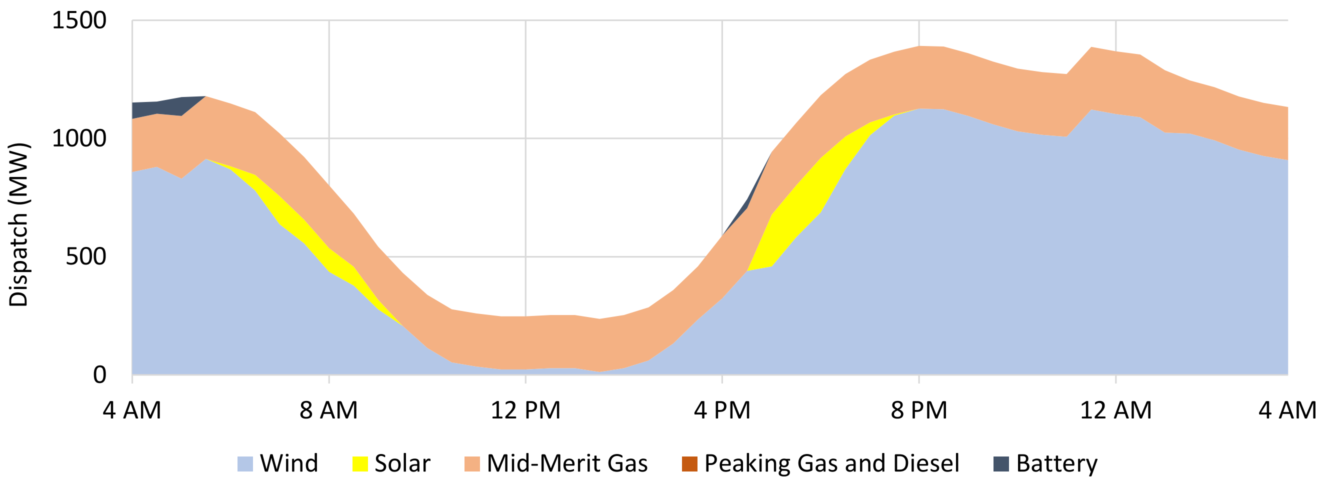

For this day, the energy price is very low in South Australia – in fact it is always negative, and is close to the market price floor of -$1,000/MWh in the middle of the day.

The three gas units provide a total of 3,122 MW.sec of inertia, and there is a further 2,200 MW.sec supplied from the two synchronous condensers. This is less than the islanded inertia requirement of 6,200 MW.sec. When the OSM is run, it must make additional commitments to fulfil this requirement. To do so, the cheapest units are chosen, which is a combination of gas units operating in synchronous condenser mode, from the Dry Creek, Pelican Point and Quarantine power stations.

While these units provide about 75% less inertia in synchronous condenser mode than they otherwise would, they are also very cheap to run (as they have no fuel cost), and hence they are selected as the most economical solution.

Figure 9 - Case 1B: Inertia schedule

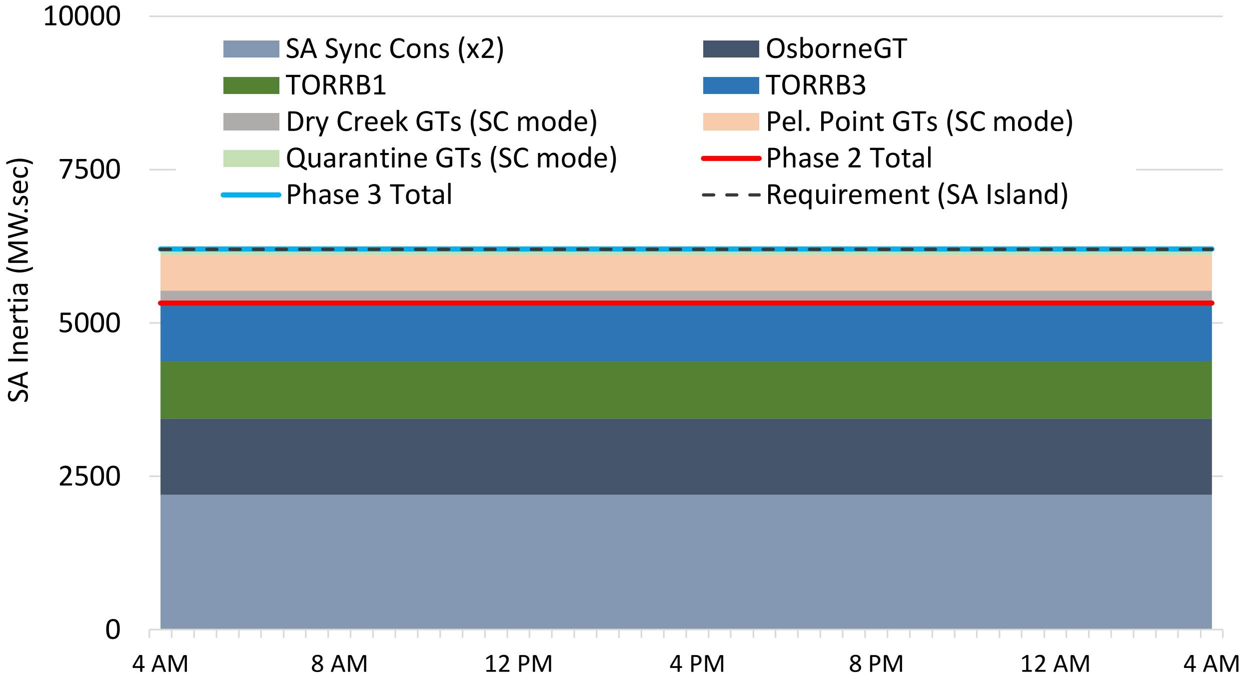

For this day, the energy price is very low in South Australia – in fact it is always negative, and is close to the market price floor of -$1,000/MWh in the middle of the day.

The three gas units provide a total of 3,122 MW.sec of inertia, and there is a further 2,200 MW.sec supplied from the two synchronous condensers. This is less than the islanded inertia requirement of 6,200 MW.sec. When the OSM is run, it must make additional commitments to fulfil this requirement. To do so, the cheapest units are chosen, which is a combination of gas units operating in synchronous condenser mode, from the Dry Creek, Pelican Point and Quarantine power stations.

While these units provide about 75% less inertia in synchronous condenser mode than they otherwise would, they are also very cheap to run (as they have no fuel cost), and hence they are selected as the most economical solution.

Figure 9 shows the inertia schedule for this schedule, resulting from the phase 3 OSM. The red line shows the total inertia that would be produced with the self-committed units online (and the SA synchronous condensers). Shown above that is the contribution from the gas units operating in synchronous condenser mode, which just exceeds the 6,200 MW.sec requirement.

Because the synchronous condensers are modelled as operating with zero energy consumption/production, there is effectively no change to the dispatch schedule or PD prices as a result of their commitment – hence these are not shown. There is an increase in dispatch costs of approximately $50,000 for South Australia, or less than $1/MWh when averaged of total SA daily demand.

Note that this outcome is entirely dependent upon the assumed costs of running the units in synchronous condenser mode (for which data was not available). These costs were deliberately assumed to be low, as there appears to be no fuel cost associated with running in this mode. Nonetheless, this example demonstrates a) a separate inertia constraint and b) dual mode unit functionality can be included within the OSM.

6.1.3 Case 1C: minimum synchronous unit constraint

The final base case tests a minimum synchronous unit constraint, and as for the previous case, uses the 27th December 2021 day as a basis. However, this case is focused on Tasmania, which is assumed not to be islanded.

For the Tasmanian region, a minimum of three large synchronous units are required for security, and for the purposes of testing this constraint, it was assumed that only five units could contribute to this constraint, being the John Butters, Gordon, Catagunya and Poatina hydro units, and the Tamar Valley combined cycle gas unit. In reality many more Tasmanian units would contribute to this constraint, such that it may be unlikely to bind in practice. Nonetheless, these assumptions allow testing of the minimum synchronous unit constraint functionality.

Outcomes are similar in concept to the previous case, and are not presented in detail. The minimum synchronous unit constraint was observed to behave as desired. In the phase 2 PD, in most intervals on two units – Poatina and Catagunya are self-committed. The OSM then commits the John Butters unit to the meet the constraint, as in Figure 10.

Interestingly, the OSM chooses to commit the John Butters unit, which was already self-committed in the morning of the test day (6am – 10am). Each Tasmanian hydro unit has the same assumed OSM enablement and running costs, and therefore each unit would have the same cost impact. However, as the John Butters unit is already starting up due to its self-commitment schedule, an *additional* start is not incurred for this unit in the OSM, making it the best choice.

Figure 10 - Case 1C: Number of online synchronous units in phase 2 and phase 3

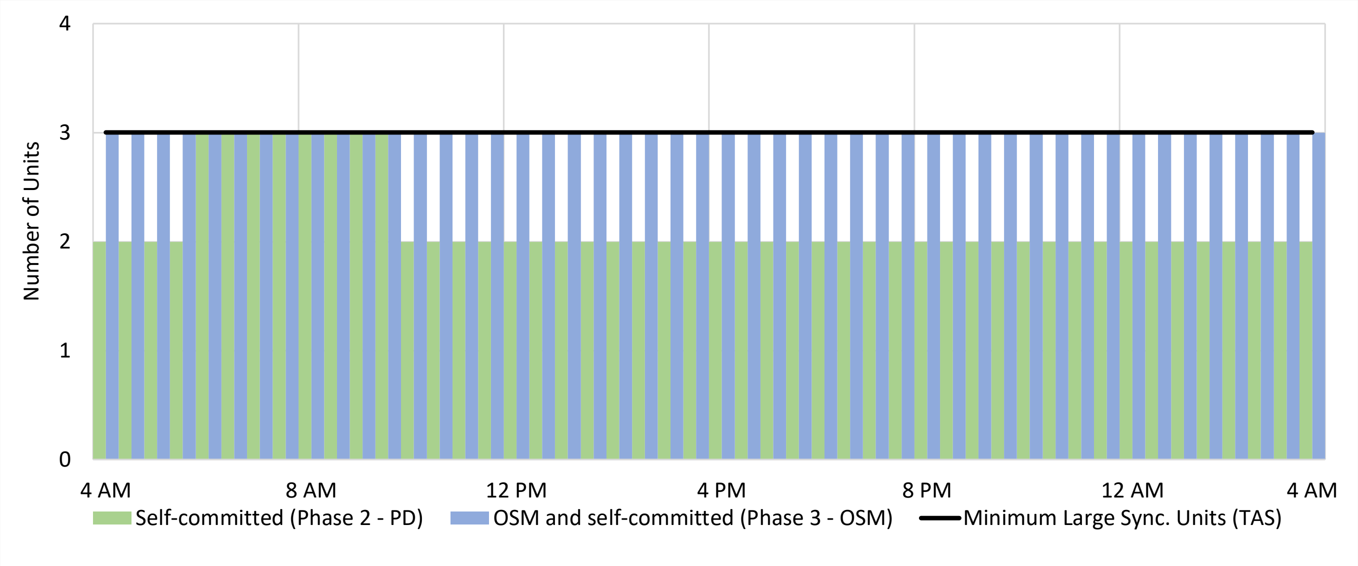

6.2 Design questions

6.2.1 Scheduling OSM-committed units in PD

By default, the pre-dispatch model in the prototype requires that units offer to submit offers that are economically dispatched to give rise to any required OSM commitments. For example, a unit that must be synchronised and running at minimum generation to be in a configuration would offer its minimum generation well below the expected energy price, likely at the price floor.

An alternative (or additional) approach is to apply constraints that require a unit to be dispatched at the required level. This is essentially what happens in the phase 3 OSM formulation itself.

When simulating scenarios using real-world data that was (relatively) unchanged, it was not observed that any units could not get scheduled at minimum generation if they submitted their energy at the market price floor. However, inputs could be deliberately altered to give rise to this outcome. For example, all wind offers could be set to -$1,000/MWh, so that the price would be at the price floor, and a necessary OSM unit would be dispatched at 0 MW for example.[20]

As such, observations are that:

- For the vast majority of time, it is likely that a participant can be scheduled in accordance with its OSM commitment simply through setting appropriate offers. This would mean constraints are unnecessary, and it has the benefit of a participant being able to 'bid out of a commitment', if it no longer intends to follow it (e.g., where the OSM commitment is not yet binding). It also has the benefit of avoiding distortions to scheduling decisions through constraints that schedule units out of merit.

- It is not impossible that a situation arises where an OSM unit cannot get scheduled. Something is required in this situation. One option could be to apply a constraint only when the OSM commitment is binding, or to allow the participant to select whether it thinks a constraint should be applied. This is an area for further consideration as policy progresses.

As a final consideration, it is noted that if batteries are scheduled in a mode to provide inertia, this is assumed to require reserving capacity to both ramp up and ramp down. Therefore, offering at the market price floor is not appropriate – the participant needs to carefully offer its capacity around the expected energy price. A challenge is that price bands cannot be updated after 12:30pm the day before the operating day. The battery operator would need to ensure its price bands are set up appropriately to give rise to a possible commitment and expected energy prices.

6.2.2 Relieving IBR constraints for net market benefits

The concept of relieving IBR constraints for net market benefits (i.e., improved gains from trade) is a fundamental aspect of the Operational Security Mechanism in general, and a particular focus area for this prototyping work. It is also a relatively novel approach to power system scheduling - this reflects some of the changes and challenges occurring in the NEM to integrate increasing quantities of inverter-based generation.

An outcome in which the OSM prototype improves the gains from trade has been tested and documented in base case 1A (Section 6.1.1). In this case, gas units were committed to allow additional South Australian wind to be dispatched, allowing market benefits to be realised via reduced thermal generation in other regions. While the improvement in market benefits was 'only' 0.5% of total costs, it is important to note that this is a small percentage of a large number.

While a simplification, conceptually it may be useful to decompose OSM commitments into two categories – those to meet minimum security requirements, and those to improve the gains from trade. Under typical offer cost scenarios (i.e., where offers reflect actual costs), the former would only increase operating costs, and would only occur where a constraint would otherwise be violated. The latter would only occur where the commitment results in increased economic efficiency, because this economic efficiency is expressed mathematically in the formulation's objective function. In reality, optimisation outcomes are complex, and a single commitment may both help to meet minimum security requirements and improve the gains from trade.

However, the key message is that the OSM formulation can identify and leverage opportunities to improve the gains from trade. Where minimum security constraints are already met by the self-committed generators, the OSM formulation would only ever make additional commitments where they result in efficiency gains.

That said, this test case has made a number of assumptions that would not hold in real scheduling outcomes, and which may mean that anticipated improvements in the gains from trade do not eventuate in real-time operation of the power system. For example:

[20] Note that the model did not include pro-rata tie-breaking that might occur in practice. This leaves the same issue – that a unit is dispatched below its minimum generation.

- Forecast errors may mean that OSM decisions are turn out to be sub-optimal with the benefit of hindsight. For example, an over-forecast of available wind may mean that OSM makes a commitment to dispatch more wind than is actually available in real-time. Ultimately this is a reflection of the challenges of decision-making under uncertainty, and not specific to the OSM formulation.

- Disorderly bidding may mean the changes in gains from trade are not accurately represented. This could occur through PD offers not accurately reflecting true costs, and perhaps also through OSM offers not reflecting true costs.

These two considerations are to be discussed further in later sections, with accompanying simulations.

6.2.3 Unbundling system services from configurations

The OSM work is being developed against the context of a rapidly evolving power system, in which engineering research is being developed on an ongoing basis. Currently, certain power system requirements are specified as being met through particular configurations of units, but other security requirements can – or in the future could be - unbundled from these configurations. Power system inertia is a good example of this, as, if engineering knowledge evolved sufficiently, it has potential to be defined as a separate requirement.

From a formulation perspective, both options – retaining bundled configurations for multiple requirements, or unbundling services where possible – are workable. That said, each configuration requires the use of additional binary variables (for each period). In comparison, an unbundled inertia constraint can be formulated using only existing commitment variables. Therefore, there may be an advantage in using an approach that reduces the number of configurations, and hence the number of binary variables. However, this advantage is only material if solve times are limiting, which then depends on other scheduling parameters.

There is also an interesting market consideration, in that it may be that unbundling services allows for each service to be priced separately, with participants submitting different offers to submit different services. However, this really depends on the nature of both the service, and the unit providing it. For example, a gas unit must be committed and online to contribute to a configuration, and if it does so, it would be synchronised and providing inertia (and energy) as a by-product. Therefore, even if it submitted a separate inertia offer component, there is effectively no physical way to clear the unit for a configuration without clearing it for inertia. In this case, having a separate cost component for inertia does not appear particularly beneficial.

6.2.4 Dual mode units

The solution to Case 1B had a number of gas units operating in synchronous condenser mode in order to satisfy a minimum inertia constraint. This has been achieved in the formulation by using separate binary variables for synchronous generation mode and synchronous condenser mode (and requiring that the unit be in no more than one of these operating modes in an interval). The formulation also allows different OSM offers and different technical properties (e.g., amount of inertia) to be associated with different operating modes.

A possible challenge associated with dual mode operation is that different units (or types of units, e.g., hydro, gas) may have different capabilities and restrictions for switching between modes, minimum off times, and so on. It is therefore not clear that a ‘one-size-fits-all’ approach would necessarily work for all unit types, but equally, maintaining several different constraint equations for different units would complicate the formulation. Future work might be to explore which constraints are most relevant or necessary for these types of units.

6.2.5 Disorderly bidding

In certain situations, NEM participants may be incentivised to engage in disorderly bidding, i.e., where their offers don’t represent the physical cost of dispatching their units. This might occur where a participant knows that its offer will not materially decrease the RRP it will be paid, but by decreasing it’s offer costs, it will enhance its ability to get dispatched.

This could influence outcomes in the OSM, because it would use these offers to calculate the gains from trade. For example, if an IBR unit offers its energy at the MPF, this will make relieving an IBR constraint appear more beneficial than if the unit offered at its true cost.

As a simple example of how this could occur, assume that:

- The “true” cost of IBR is $0/MWh, and the marginal cost of energy that would be displaced cost is $100/MWh.
- An OSM unit can be committed with an enablement cost of $100,000 and a variable cost of $100/MWh (so that its minimum generation neither increases nor decreases dispatch costs, and only the enablement cost is relevant).
- Committing this unit allows for an additional 200 MW of IBR, saving $20,000/hour, which is the improvement in hourly gains from trade if the IBR is offered at its actual cost.

If this scenario occurs for only four hours, it is not worth committing the unit, as the benefits of $80,000 are less than the enablement cost. If, however the IBR is offered at -$1,000/MWh, the benefit becomes $220,000/hr and the unit is worth committing, even for a single hour. This scenario (which is depicted in Figure 11) is intentionally simplistic, but the same outcome can be replicated with the prototype.

Figure 11 - Increase in apparent gains from trade when IBR is offered below actual cost

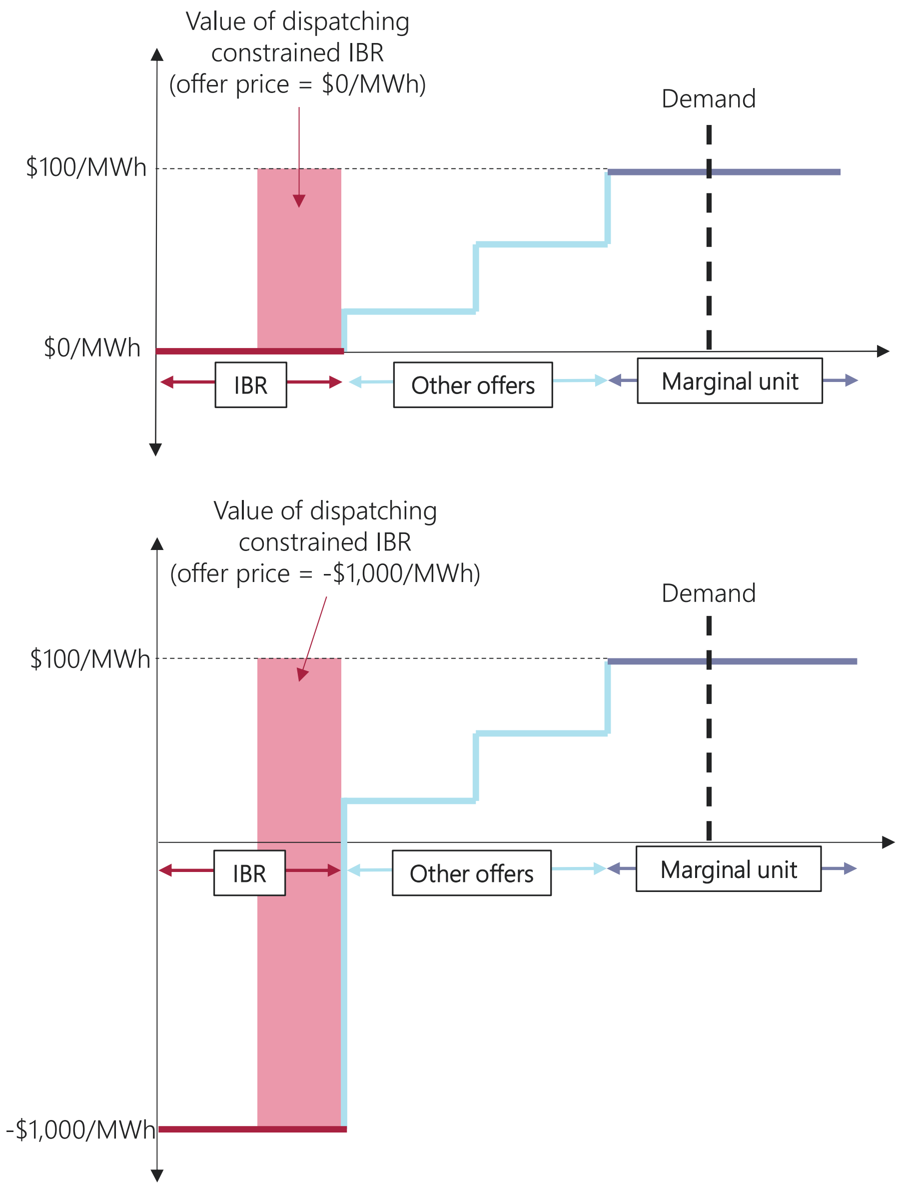

Two main factors contribute to this outcome, being that a) the cost of committing the OSM units is not recovered entirely/directly from the IBR that is relieved and b) the zonal structure creates scenarios where offering below cost does not materially affect the price a participant is paid. Of course, the latter can create disorderly bidding incentives with or without the OSM.

This adds a level of dynamism to the OSM that may make participation challenging. It may be relatively straight-forward for a participant to predict whether it will be OSM committed for a static security constraint. It would be more difficult to predict whether it will be committed to improve the gains from trade when this depends on (disorderly) bidding outcomes. This may be particularly problematic for units that have a notification time longer than the cut-off time, where they may or may not be needed if the IBR changes its bids.

6.2.6 Potential to improve the gains from trade

The ability for the OSM to make commitments that indirectly improve the gains from trade by relieving constrained IBR is novel in the context of scheduling approaches in international power systems. In principle, the model is able to appropriately make these decisions, for example trading off the additional cost of committing more expensive units where this may be more than offset by additional low-cost wind or solar generation. In practice, there may be limited ability for this outcome to occur, given the nature of the current understanding of these IBR limits.

The potential for improvements to the gains from trade are dependent upon the potential to find a configuration that allows IBR generation, above that associated with the minimum-security constraints. Currently, there are 95 published configurations for SA when in a normal (non-islanded state), which can be considered across three categories based on whether they require either zero, two or four SA synchronous condensers operating (Table 16).

Table 16 - Maximum IBR in currently-defined configurations

NUMBER OF SA SYNCHRONOUS CONDENSERS	MAXIMUM GENERATION FROM IBR	COUNT OF CONFIGURATIONS
NORMAL SA OPERATION		
4 synchronous condensers	2500	21
2 synchronous condensers	2000	8
	1900	16
No synchronous condensers	1750	1
	1700	12
	1650	1
	1600	4
	1550	1
	1450	10
	1400	5
	1350	5
	1300	11
ISLANDED SA OPERATION		
2 synchronous condensers	1900	11
	1800	4
0 synchronous condensers	1300	38

7 SCHEDULING PARAMETER TESTS

7.1 Overview

This section explores the impact of various combinations of so-called scheduling parameters, such as block duration and granularity of enablements.

7.1.1 Scheduling parameters

Table 17 sets out the scheduling parameters for testing during the prototype development, with default and alternate values. Figure 12 provides a conceptual view of these scheduling parameters.

Table 17 - Scheduling parameters for testing

SCHEDULING PARAMETER	DEFAULT CONDITION	TEST CONDITIONS	COMMENT
BLOCK SIZE	4 hours	8 hours, 2 hours	
GRANULARITY OF ENABLEMENTS	2 hours	8 and 4 hours, 30 mins	8 hours was also tested.
TREATMENT OF HORIZON IN OPTIMISATION	Optimise over multiple blocks	Treat blocks sequentially	Optimising over a single block can produce end effects.
TIME BETWEEN CUT-OFF AND BLOCK START TIME	2 hours	1 hour, 4 hours	Cut-off time was modelled as affecting forecast (wind/solar/demand) quality. Note that real forecasts were not used, although forecasts errors were determined with referenced to observed error magnitude. This made the exact value of the cut-off timing somewhat arbitrary.
TIME BETWEEN GATE CLOSURE AND BLOCK START TIME	4 hours	2 hours	Algorithm run time may constrain this parameter. Timing of gate closure with respect to PD timetable will determine the number of re-bidding iterations. See comments below.

Other optional parameters that were considered for testing were the granularity of inputs (e.g., forecasts, PD offers, constraints), and the OSM horizon, but ultimately these parameters are left as potential future work. That said, it seems sensible to use an input granularity that is aligned with the pre-dispatch (30-minutes). There may indeed be value in testing a longer OSM horizon in the future.

Figure 12 – Conceptual view of scheduling parameters

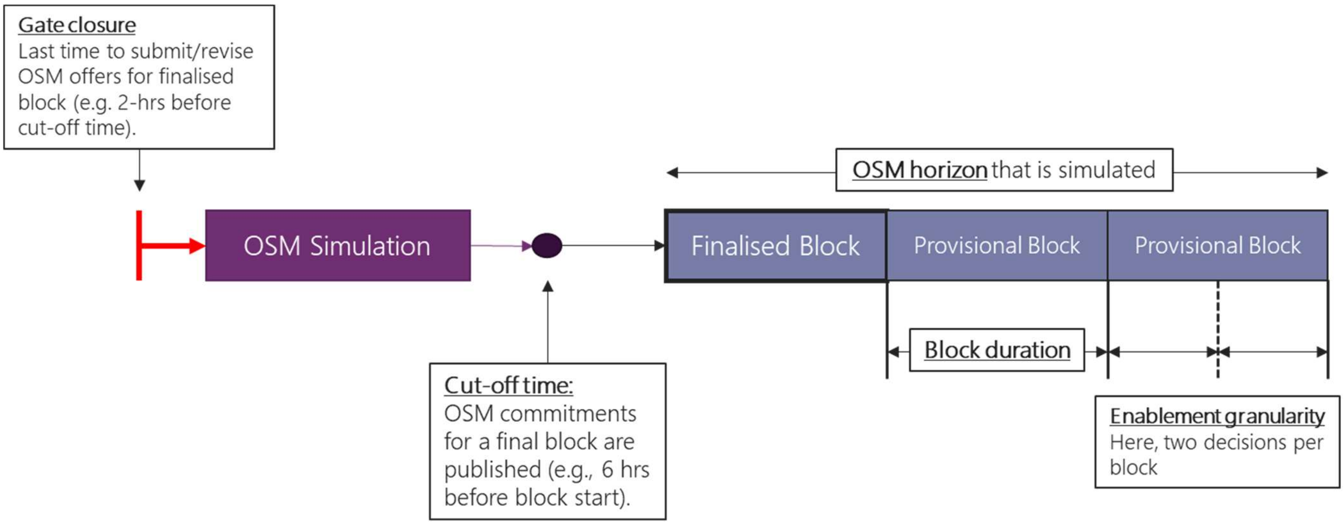

7.1.2 Performance metrics

Various metrics are considered in the assessment.

- Solution time of the OSM process (i.e., phase 3 only).
- Economic efficiency, e.g., gains from trade.
- Shortfalls in meeting security needs.
- Predictability as measured by the variability in the choice of units – and associated costs – scheduled for OSM across the day.

The following will be more holistic:

- Projected solution time (for a high-end commercial solver or for servers with large number of threads). Drawing on our experience of different solvers, approximate estimates of how performance might vary with higher and lower performance solvers.
- Software footprint (i.e., hardware needed/assumed for a commercial solver)

7.2 Gains from trade assessment

7.2.1 Block size and cut-off

Figure 13 shows the improvement in gains from trade in phase 4 relative to phase 2 (i.e., before and after the OSM commitments). This is shown for a variety of cut-off times (CO), and block durations (B), and also for two different ‘directions’ of forecast errors. Note that a longer cut-off time means decisions must be made with more forecast error, while a longer block duration means decisions must be made further ahead.

We first note that the change in gains from trade is improved by between $3-4M over the operating day, and differences between the improvement in gains from trade for different scheduling parameters is less than $1M. This suggests the decisions being made are not hugely sensitive to the choice of scheduling parameters.

The first three bars show outcomes for increasing block size with constant cut-off of 1 hour. With increasing block size, the gains from trade improvement is slightly reduced in the high net demand case (and near-constant in the low net demand case). This outcome occurs as longer block size means decisions are being locked in further ahead of time, i.e., with more uncertainty. Increasing cut-off time similarly decreases the gains from trade, as there is greater forecast error.

Figure 13 - Change in gains of trade from phase 4 to 2 (optimisation over all blocks in 24-hour period)

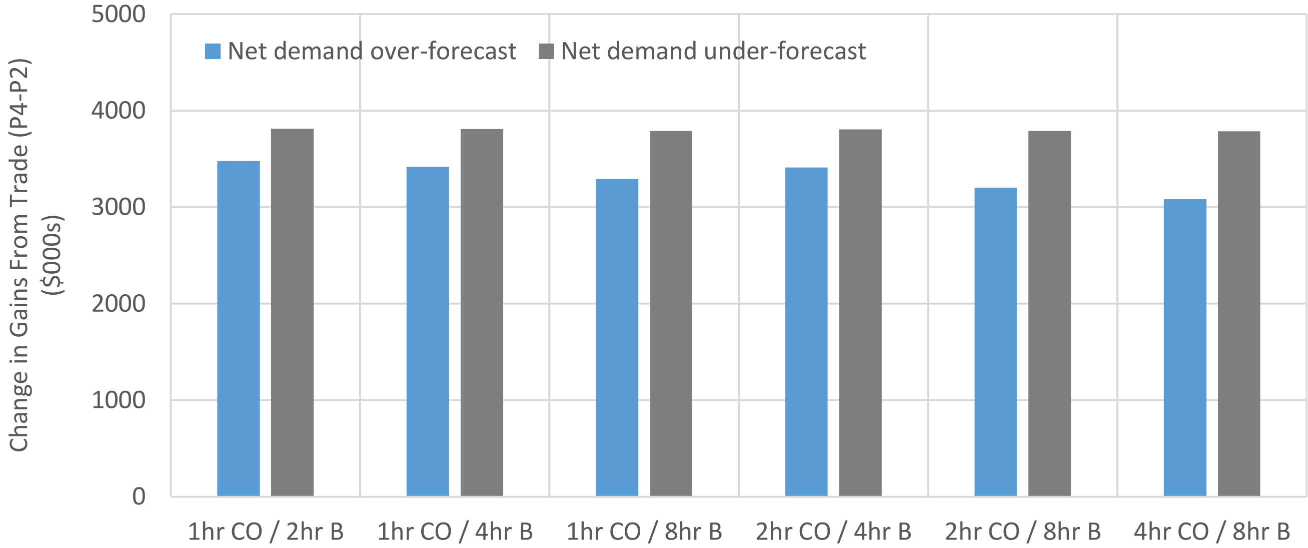

While there are variations in solution quality between different scheduling parameters, these are small compared to differences between simulations using different forecasts. Based on the gains from trade, there does not appear to be a particular combination that stands out as being preferred, particularly for block size where both 2-hours and 4-hours give similar outcomes.

7.2.2 Single block optimisation

Figure 14 shows similar results, but with the optimisation over a single block only. For example, with an eight-block duration, decisions are made with information available only for those eight hours. In contrast, results in Figure 13 are made with a full 24-hours of information, but with results locked in for the first 8-hours only. The second block of eight-hours is then optimised with the 24-hour horizon rolled forward 24-hours (and so on for the third block).

It is first noted that the improvement in gains from trade are not as good as for the previous section, because there is no information about what will happen after a block ends, and there are likely to be end-effects. For similar reasons, it is better to have a longer block duration, as this means there is more information included in the optimisation.

It does not appear advisable to use a single block approach with a block duration of less than 8 hours. In fact, from an optimisation point of view, the single approach will be at best as good as optimising over all blocks. As such, the single block approach is not recommended unless there are strong policy reasons that support it.

Figure 14 - Change in gains of trade from phase 4 to 2 (optimisation over first block only)

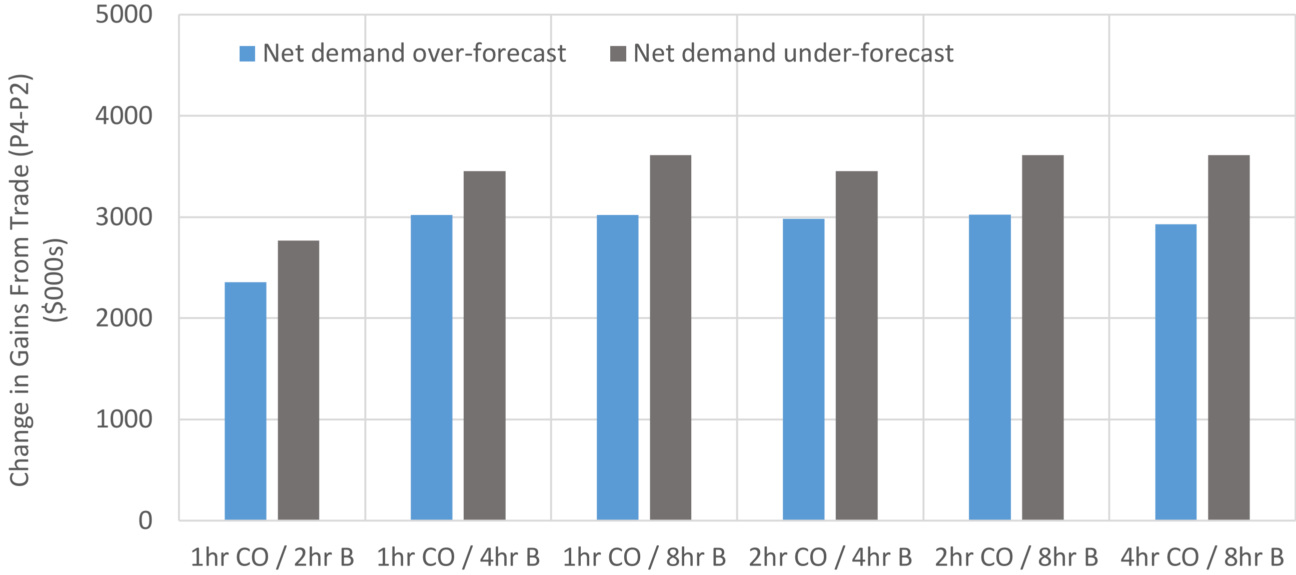

7.2.3 Granularity of commitments

This section compares performance of a range of granularity of enablement durations, herein referred to as the commitment interval duration. The tested values were 30-minutes, 1 hour, 2 hours, 4 hours and 8 hours. All simulations use a 4-hour cut-off time and 8-hour block. The case with 8-hour granularity can be considered equivalent to a 'fixed' scheduling approach, in which each unit is either on or off for an entire block (block duration is the same as commitment granularity). This means that there is no ability for the solver to trade-off between units with either high start-up costs and low variable costs, or vice versa, because any commitment must run for the full block. However, the distinction between start-up and variable costs may still come into play if it is needed or economical to run a unit for more than one block.

In principle, a longer commitment duration gives less flexibility. For example, if it were precisely optimal to commit a unit from for 1.5 hours across the period 12.30pm to 2pm, then

- With 30-min granularity, this could be exactly achieved.
- With 2-hour granularity, the unit must be committed from 12pm to 2pm, and with 4-hour granularity, from 12pm to 4 pm (note that the first interval begins at 4am).

However, outcomes also depend on other factors, such as a unit's minimum run time, or the ratio of its enablement and running cost (it may be better to keep a unit with a high enablement cost on for a few intermediate periods if it is needed again later).

On the other hand, a longer commitment interval might give simpler, more predictable outcomes (though these are difficult to measure), and faster solve times. The fixed scheduling approach (block length is the same as commitment granularity) may also make it simpler for participants to determine how to structure their OSM offers, because there they are essentially submitting a single cost to run for the entire block.

Table 18 shows gains from trade in phase 4, as well as the percentage different relative to the 30-minute case (first row).

Table 18 - Comparison of phase 4 gains from trade under varying commitment granularity

GRANULARITY	NET-DEMAND OVER-FORECAST		NET-DEMAND UNDER-FORECAST	
	GAINS FROM TRADE ($M)	% CHANGE (REL. TO 30-MIN)	GAINS FROM TRADE ($M)	% CHANGE (REL. TO 30-MIN)
30-MIN	$769	-	$768	-
1 HOURS	$768	-0.1%	$767	-0.1%
2 HOURS	$767	-0.3%	$766	-0.3%
4 HOURS	$763	-0.8%	$762	-0.9%
8 HOURS	$758	-1.5%	$757	-1.4%

Outcomes are broadly similar with both directions of net-demand forecast errors, where the gains from trade are lower with longer granularity.

There is little benefit in the gains from trade when using a 30-minute or even 1-hour duration compared to 2 or 4 hours. This is not surprising given that most units were assumed to have a minimum on time of one-hour, and the key OSM units (e.g., Torrens Island, Pelican Point, Osborne) had minimum on times of four hours.[21] In Victoria, it is the Loy Yang (A and B) and Yallourn units that are most important (all configurations require at least three of these units), but they have minimum on times of 16 hours.[22]

Comparing the OSM commitment schedules with 30-minute and 8-hour commitment interval duration for South Australia and Victoria show that units from Hallet, Ladbroke, and Quarantine (in SA) and Bairnsdale, Mortlake and West Kiewa (in VIC) are committed for longer durations in with 8-hour commitments. These units aren't strictly required for the OSM security constraints, and it appears that they are committed for energy for durations less than 8-hours at 30-minute resolution. Therefore, it may be that if the proposed approach to avoid non-security commitments were used, these differences would be removed.

In conclusion, given that the key units for the currently defined configurations have minimum on times of four hours or more, it does not appear that there is much value in using a configuration of less than 2-4 hours.

7.3 Predictability and certainty

Predictability (how much variation there could be between runs) is challenging to thoroughly assess, but overall, the prototype results appear acceptable in this respect.

Predictability could be low in two key respects:

- **Re-bidding:** A new OSM commitment impacts prices and therefore a self-committed unit de-commits, requiring a new OSM commitment in order to meet all security requirements (which could further depress prices!)
- **Similar solutions:** Subsequent OSM runs (with updated forecasts etc) choose different units – e.g., there are many solutions with similar costs, such that a small change can 'flip' to another solution.

It was found that OSM commitments for security and relieved IBR constraints were generally very constant across the day – e.g., the same units were committed, without needing to constantly start or stop units. This may be partially due to the existence of start-up costs themselves, and also because the OSM variable costs were constant

[21] These values were obtained from the 2022 ISP.
[22] These values were obtained from the ACIL Allen Fuel and Technology Cost Review, 2014.

by period (which is understood the policy intent). Therefore, the solver would normally not swap between different OSM commitments across the day. The solver was also observed to choose a unit that was planned to be self-committed at some point in the day, as this allowed for an avoided start-up cost (though this will not necessarily be a general result).

It could be valuable to undertake analysis of historical directions events, to understand how often and for long such events have occurred. For our selected days, OSM commitments were typically required for all or most of the day, which contributed to predictable solutions, even with short commitment granularity.

Furthermore, the configurations themselves appear to provide for solution predictability and certainty. For example, in SA it is sufficient to have two SA synchronous condensers and two thermal units online (e.g., two Torrens Island B units or one Torrens Island B with one Osborne). It follows that at most two iterations of OSM and PD re-bidding would be required to meet minimum security constraints. In practice, it could take longer if participants engage in involved re-bidding and updates to self-commitment flags, but such outcomes cannot practically be modelled.

It was also found that there were limited changes to OSM commitments in order to relieve more IBR. This is because most configurations allow similar amounts of IBR (e.g., with two synchronous condensers online, all configurations allow between 1,900 and 2,000 MW). These limits are relatively high compared to IBR capacity in SA, and it was found to require particular combinations of demand and IBR availability to warrant incurring additional costs to reach a configuration that allows for 2,000 MW. Of course, these conclusions may change as more IBR is installed in SA, or as there is greater ability to export IBR to other regions such as via Project Energy Connect.

In some simulations, it was found that synchronous condensers could be cycled on and off because they had very similar or identical costs. This was an artefact of there not being available cost data for these units, and would not be expected to occur in reality (unless perhaps many units offer at zero cost).

7.4 Solution time and software footprint

7.4.1 Prototype and production run time

Run time varied from approximately 15 seconds (single block with 2-hour block duration and 1-hour commitment interval) up to 14 minutes (all blocks over 24-hours and 1-hour granularity). In general, the fastest runs were those that were over a single block, which all required less than a minute, while for other scheduling parameters the inputs (demand, wind/solar forecasts, PD offers etc) themselves affected run time more than the choice of scheduling parameter. Average run time was less than two minutes.

Interestingly, commitment duration only weakly affected solution time, with 30-minute decisions requiring approximately 3 minutes, and other durations requiring approximately 2-minutes.

Note that for an optimisation over 24-hours (i.e., all blocks), block duration would not affect run time as block duration only affects which of the OSM commitments are locked in.

These solution times apply for the FICO-Xpress solver on a standard desktop PC. It is noted that there was material variation between solve time for different days, with the 2nd August 2022 day requiring – on average - approximately half as long as the 27th December 2021 day. Such outcomes are not unexpected, with the topography of the problem affecting how quickly the solver arrives at an optimal solution.

While these results provide some indication of performance, they do not necessarily imply anything concrete about solve times for a production version with these varying due to:

- Possibly requiring all the generic constraints that are included in NEMDE,
- Requiring additional binary variables for interconnectors,

- Being able to use higher-end hardware,
- Being able to tailor the solution algorithm for better performance,
- Having more or different security constraints and configurations,
- Having either more or less units participating in the OSM than was assumed here.

7.4.2 Solution footprint

Scheduling parameters do not have any real impact on solution footprint – it is the core formulation itself which sets solver requirements.

The developed formulation is a MILP with similarity to generic unit commitment problems. The UC problem is classified as NP-hard, so that exact solutions cannot be found for large problem sizes, though approximately optimal solutions can be found.

The prototype was tested with both a high-end commercial solver (FICO-Xpress), and an open-source solver (CBC). It was found that that many of the cases could not be solved in a reasonable (<1 hour) amount of time using the latter. While there are alternatives to FICO-Xpress that offer good performance, some form of performant commercial solver would be required to run the formulation for the real NEM.

The prototype was written in GAMS (a commercial optimisation language), but alternative optimisation packages (of which there are many) could be used.

8 SUMMARY OF FINDINGS

This section compiles various outcomes, learnings and recommendations observed in the course of developing the formulation and exploring case studies and scheduling parameters.

8.1 Summary of findings

8.1.1 OSM functionality

Findings relating to the basic OSM functionality are:

- **Configurations and other security constraints:** The basic functionality of resolving different types of security constraints has been demonstrated, with the model being used to commit units to achieve a secure configuration, as well as for inertia or a minimum number of synchronous units. Other types of constraints could in principle be included if they can be formulated as linear, or piecewise linear, combinations of relevant variables, though of course this may increase solution times.

- **Selecting a configuration to improve the gains from trade:** The OSM was also used to demonstrate a case in which it makes a commitment to improve the gains from trade. This is reflected in a subsequent PD run by changing the RHS of a constraint (for example on allowable IBR in a region). In general, any RHS parameter could be updated as a result of decisions made in the OSM, so long as the costs or benefits of changing that parameter are included in the OSM formulation.

- **Dual operating modes:** The formulation is capable of incorporating units that can operate in more than one mode, with different cost components and technical attributes in each mode, as well as batteries operating in a mode in which they provide inertia. The assumptions used for battery inertia would need to reviewed should this functionality be incorporated.

- **Energy-only commitments:** OSM policy is that the OSM should not make commitments where there is no additional security benefit, e.g., where the unit is committed because its energy is lower cost than units being scheduled through the PD. A two-step OSM run has been proposed to give effect to this, but not yet implemented within the prototype.

8.1.2 Solver performance and solution footprint

Findings relating to the formulation and OSM performance are:

- Solution time: Solve time for the OSM was observed to have a strong dependence upon the number of units participating in the OSM. This is not unexpected given that each OSM unit requires a number of binary variables (for commitments, start-up, shut-down, etc) per commitment period. That said, if an OSM unit is self-committed, then this ‘fixes’ its binary variables for that period, and this would tend to reduce solution times. On average, the OSM solution time is approximately 3 minutes, and with the longest solution time being under ~15 minutes, for a case with very low demand. For reference, the XPRESS solver with an optimality tolerance (or gap) of 0.01% was used, on a typical desktop PC.

- Production solve time: There is uncertainty in extrapolating to estimate solve times for a production version of the OSM. This depends on many factors, including the impact of a full set of constraints (if required), and possible reductions in solve time from using more performant hardware. The main identified requirement is that a commercial solver would be required.

- Solver and other requirements: The model was also tested with an open-source solver (CBC), but this did not always obtain a solution. A reasonably high-end commercial solver should be used for the OSM. The model was written using the GAMS language, but this is not restrictive and another appropriate optimisation module could be used.

8.1.3 Scheduling parameters

Findings related to the scheduling parameter testing are:

- Cut-off time and block size: A range of cut-off times (which increases forecast uncertainty at the time decisions are made) and block size (which increases the horizon over which decisions are finalised) were tested.
  - When optimising over all blocks in a 24-hour horizon (but finalising decisions for the first block only) these parameters had relatively low impact on the gains from trade. In principle a small block size and short cut-off time allow for the best solutions, but the potential benefits may be small enough that there is greater value in being able to finalise decisions earlier, giving more certainty for the operator and participants.
  - When optimising over a single block only, a longer block size is preferred, with eight hours likely to be a minimum size. This approach is not recommended unless there are strong policy reasons to do so, as the solution will not take into account any information beyond the end of the block.
- Time resolution: Commitment granularity was tested for durations of 30-minutes, and 1, 2, 4 and 8 hours. While a shorter time gives the best decisions (because there is more flexibility in when units are committed), these benefits were not material below 2-hours. Equally, a longer duration makes decisions simpler and more predictable. There may be rationale in adopting a duration of 4 or 8 hours - at the cost of some economic efficiency. Note that several key security units were steam turbines with a minimum on time of at least four hours with the developed datasets, which reduces the value of making commitments at 30-minute or 1-hour granularity.
- Predictability: Issues over predictability of OSM commitments did not generally arise, except in cases where many OSM offers were structured to be similar (e.g., for synchronous condensers were assumed to have low costs). For most simulations, it was optimal to commit one or two units, and run them constantly over the day, as this avoids unnecessary start costs. We note that predictability is difficult to test without knowing how participants would construct their OSM offers or set their self-commitment flags based on the spot market.
- Impact of scheduling parameters on solve time: Solution time was not strongly affected by most scheduling parameters, with the exception of a single block optimisation which significantly reduces the number of decisions which need to be made, and never required more than one minute to solve. On average, the OSM run time (for a single solve) was approximately 3-minutes, but in one case was up to 14 minutes. It was found that inputs such as demand, wind/solar availability, self-commitment flag settings and offers could have greater impact on solve times than the scheduling parameters.

8.1.4 OSM design

Findings related to the overall performance of the OSM design are:

- Separate and/or unbundled system services: The modelling tested an inertia constraint that was specified separately to the configurations. The model did not use separate offers for this service, as those costs were already reflected in the start-up and running costs used for the configurations. It would be possible to test an unbundled service (with separate offers) with the prototype, but how this would work depends on the nature of the service. A challenge for inertia is that it is provided by synchronous units as a direct result of being synchronised, which is the same requirement to be chosen for a configuration. It would be useful to develop a design for this type of unbundled service before proceeding with prototyping of unbundled security service markets.
- Potential to improve gains from trade: The potential for improvements to the gains from trade are dependent upon the potential to find a configuration that allows IBR generation, above that associated with the minimum-security constraints. However, with the currently-defined configurations the major decision that affects the IBR limits is the number of SA synchronous condensers – which is not included within the OSM decision-making scope.

- Approach to scheduling OSM units in PD: OSM units could be scheduled through updating offers accordingly, or by being constrained on. While it is logical that participants submit offers consistent with their OSM commitment (e.g., at the price floor if they must run at their minimum generation), in some cases this may not be sufficient, for example, if the price is being set at the floor because many non-OSM units are also offering at the floor price. This was observed as a possible outcome in our version of the NEM PD. This did require specifically set-up input data (e.g., all wind capacity offered at -$1,000/MWh), so this outcome may not commonly occur in practice - but in principle it could occur. While it may not be necessary to use constraints in PD most of the time, it is recommended that there be some mechanism by which an OSM unit can ensure it will be scheduled beyond its PD offers themselves.

8.2 Recommendations

As a result of this work, a number of recommendations are made and summarised below.

Scheduling parameters:

- It is suggested that a single block optimisation not be progressed, and optimisation occur over at least 24-hours, even if decisions are only being finalised for a subset of that time. While not strictly necessary, there may be logic in using 48-hours is this computationally tractable.

- There does not appear to be significant additional value in using a commitment interval duration of less than 2-hours; while the solutions are theoretically better with a shorter duration, 2-hours is found to be the point of diminishing returns. Using a longer duration gives more certainty and simpler outcomes which may suggest 4-hour or longer durations may be preferable.

- Block duration and cut-off have relatively low impact on the gains from trade. A longer block duration and a longer cut-off both increase uncertainty, and hence may have less optimal solutions, but this appears small and there may be other concerns (e.g., operator and participant certainty) that dominate these decisions. This may particularly be the case for units with a long notification time.

- In sum, it is found that there are multiple combinations of scheduling parameters that could work, for example both 4-hour and 8-hour blocks gave similar outcomes. Therefore, other implications of these parameters should be canvassed with e.g., control room staff/operators and stakeholders as preferred choices of these parameters will also depend on considerations that cannot be modelled in the prototype.

Design questions:

- Further work should be undertaken both to define the situations in which it is or is not acceptable to commit a unit through the OSM, and to test and assess the proposed approaches to avoiding energy-only commitments. It may be that the proposed approach does not guarantee the policy intent would be achieved in some (perhaps unlikely but not impossible) scenarios, and it should be considered the extent to which this is acceptable. It may also be worth considering this in conjunction with incentives to offer directly into the regular energy/FCAS through OSM settlement.

- It was not observed that the real-world data resulted in cases where an OSM unit could not get scheduled in the subsequent PD simply through submitting appropriate PD offers. However, this situation could arise where data was specifically modified to do so. It is recommended that there be a process available to ensure OSM committed units can get scheduled when the price is at the market price floor, though this may be a very uncommon outcome.

- Disorderly bidding was observed to result in outcomes where an OSM unit would be scheduled, where that unit would not be scheduled if PD bids represented actual costs. Incentives to do this may arise due in-part to the NEM's regional structure, but also that OSM costs may not be recovered directly from those who benefit. Such outcomes should be considered in progressing OSM policy.

Assessment of frequency of OSM commitments:

- This work focussed on modelling of days that were identified as strong candidates for requiring OSM commitments, e.g., where there would be likely to be no/few synchronous units online, or where IBR is likely to be constrained. These are likely to overstate the value of the OSM on a typical day.
- It would therefore be valuable to assess how frequently OSM commitments might need to made, based on either historical use of directions, or forecasts of future outcomes.
- This could also give an idea of how long security requirements might typically be binding (e.g., all day, only around midday, etc), which in turn will inform scheduling parameter decisions.

Formulation implementation:

- While the prototype provides a useful basis for planning for possible OSM implementation, it is important to note that it includes a lot of functionality for testing various approaches to the OSM that would not be needed in practice. Once the various scheduling parameters are decided, it could be possible to simplify some aspects of the formulation. Additionally, there are often multiple ways to formulate optimisation problems, and it should not be assumed that the prototype formulation is the only possible approach.
- Consideration should be given to how alignment is needed between the constraints in OSM, and in other processes e.g., pre-dispatch and NEMDE. This is particularly relevant for transmission generic constraints, loss factors (particularly on interconnectors). One would expect that including all the detailed constraints that are included in NEMDE would increase run times. Equally, it may not be necessary to include all these constraints in order to get a reasonable solution.
- It is also cautioned that the prototype and production solve times may differ materially. The solve times for the prototype were not prohibitive, but this does not necessarily mean solve time in production is no longer a concern.
- The formulation made some assumptions about operation of dual mode units, such as whether a delay is needed between switching from normal operation to synchronous condenser operation. A challenge is that it is that a one-size-fits-all approach may not work for all units. It may be necessary to include a variety of different constraints for different units, or simply to require participants to manage these constraints internally.
- For the purposes of testing, assumptions were also made regarding batteries providing inertia, and these should be reviewed as more information becomes available.

APPENDIX A MATHEMATICAL FORMULATION

This appendix presents the mathematical formulation, and is current as of version 1.3 of the GAMS model.

A.1 Sets

Table 19 describes the sets used in the OSM model formulation.

Table 19 - Sets in the OSM formulation

SET	DESCRIPTION
$d \in D$	Dispatch intervals (dz is an alternative index used for data import only).
$c \in C$	Capacity (or commitment) intervals
$r \in R$	Regions
$sr \in R_S$	Sub-regions
$f \in F$	FCAS service types [6 second, 60 second, 5-minute, regulation][23]
$fx \in F$	Contingency FCAS service types [6 second,60 second,5 minute]
$fd \in F_d$	FCAS directions [raise, lower]
$t \in T$	Transmission links
tLnkO(tLnk, sr)	Origin sub-region for the transmission link
tLnkD(tLnk, sr)	Destination sub-region for the transmission link
$gc \in G_C$	Generating units capable of operating in synchronous condenser mode
$gsr \in G_{SR}$	Generating units in sub-region sr
$gr \in G_R$	Generating units in region r
$gs \in G_S$	Set of units that contribute to minimum synchronous units constraint.
$gi \in G_{IBR}$	Set of IBR units (wind and solar).
$csg \in CSG$	Groups of units that can exist in a configuration.
$gx \in G_X$	Storage units
$gb \in G_B$	Battery units
$gtype \in GT$	Generator types
$so_{SA} \in SO_{SA}$	Security configurations for South Australia
$so_{VIC} \in SO_{VIC}$	Security configurations for Victoria
$s \in S$	PDS energy and FCAS bid and offer steps [1, 2, … 10]
$ccmp \in CCMP$	Cost elements for reporting, includes the following: PDS_Bids, OSM_Running, OSM_Startup, OSM_Startup_Discount, OSM_PostBids, FCAS, Structural, SynCon, Infeasibility

[23] The model does not include the Fast Frequency Response (FFR) services to be introduced in October 2023, but in principle this could be included similarly to the other FCAS types.

A.2 Parameters

Table 20 sets out the various power system and regional parameters (i.e., fixed values/constants) used in the prototype. Table 21 and Table 22 set out parameters relating to generating units and other resources. Table 23 lists the scalars.

Table 20 - Power system and region data

PARAMETER	SYMBOL	SETS	UNITS	COMMENTS
DEMAND AND SECURITY REQUIREMENTS				
Demand by subregion and dispatch interval	*rDmd*	sr, d	MW	
Local FCAS requirement for subregion	*MnFCAS_SR*	sr, d, f, f_d	MW	Only takes effect in islanding scenarios
Flag: Local FCAS requirement applies for subregion	*MnFCAS_SRF*	sr, d, f, f_d		
Global FCAS requirement	*MnFCAS_Glb*	d, f, f_d	MW	
Inertia requirement for region (secure operating level of inertia	*rInertia*		MWs	
Required number of online synchronous units in region	*rSUnits*	r	# units	
Required number of online synchronous condenser units in region r	rSUnits_SynCon	R	# units	
TRANSMISSION LIMITS				
Transmission capacity (forward direction)	*Trans_CapF*	t	MW	
Transmission capacity (reverse direction)	*Trans_CapR*	t	MW	
Transmission loss factor	*Trans_Loss*	t	frac	Set to 0 now.
SECURITY CONFIGURATION LIMITS				
Maximum IBR generation in SA under security configuration	*SCnst_MxWS*	so_SA	MW	
Maximum imports in Victoria under security configuration	*SCnst_MxImp*	so_VIC	MW	
Minimum number of units required for security option and configuration security group (SA)	*SCnst_MinUnits_SA*	so_SA	# units	
Minimum number of units required for security option and configuration security group (VIC)	*SCnst_MinUnits_VIC*	so_VIC	# units	

Table 21 - Resource parameters

PARAMETER	SYMBOL	SETS	UNITS	COMMENTS
GENERATING UNIT PARAMETERS				
Capacity	*GenMxGen*	g	MW	
Minimum stable generation	*GenMnGen*	g	MW	Hydro-thermal units only
Inertia	*GenInertia*	g	MWs	
Minimum on time	*GenMnOn*	g		Specified as number of dispatch intervals
Minimum off time	*GenMnOff*	g		Specified as number of dispatch intervals
Maximum ramp up rate	*GenRmpUp*	g	MW	Specified as MW per dispatch interval.
Maximum ramp down rate	*GenRmpDown*	g	MW	Specified as MW per dispatch interval.
Forecast variable resource availability	*GenWindSolarF*	g, d	MW	
Flag set to 1 if unit contributes or can change mode to contribute to minimum synchronous units constraint (0 otherwise)	*Gen_SynFlg*	g		
Fraction of Unit Inertia available if it is run in synchronous mode.	*Gen_SynCvnIS*	gc	Frac	
Number of dispatch intervals delay for unit before switching from generation mode to synchronous mode	*Gen_SynCvnDly*	gc	#dispatch intervals	
	Gen_SynConHCst	gc	$	
Mapping from units (g) to	*gMapcsg*	g, csg	NA	1 if unit g is a member of constraint station group csg, 0 otherwise.

PARAMETER	SYMBOL	SETS	UNITS	COMMENTS
constraint station groups (csg)				
ADDITIONAL PARAMETERS FOR STORAGE UNITS (BATTERIES, PUMPED HYDRO)				
Energy storage capacity	*GenMxStorVol*	gx	MWh	
Storage injection rate (before losses)	*GenMxInjRte*	gx	MW	
Storage injection losses	GenInjLoss	gx	Frac	
Withdrawal losses	*GenWthLoss*	gx	Frac	
FCAS PARAMETERS[24]				
Capability for FCAS service	*FCAS_Cap*	g, f, f_d	MW	Zero for all units that do not participate in FCAS markets
OFFER DATA (PDS AND OSM)				
PDS energy offer step quantity	*PDS_BidQ*	g, d, s	MW	
PDS energy offer step price	*PDS_BidP*	g, d, s	$000/MWh	
PDS FCAS offer step quantity	*FCAS_BidQ*	g, f, f_d, d, s	MW	
PDS FCAS offer step price	*FCAS_BidP*	g, f, f_d, d, s	$000/MW	
PDS energy offer step quantity used to calculate *PDS_BidQD*	*PDS_BidQC*	g, d, s	MW	Used to build PDS_BidQS and contains cumulative bid quantities. For OSM units this is cumulative bid quantities minue minimum generation for the units. For PDS schedule units this is cumulative PDS bid quantities.
PDS energy offer step quantity used to calculate *PDS_BidQS*	*PDS_BidQD*	g, d, s	MW	Used to build PDS_BidQS and contains cumulative bid quantities. For OSM units this is cumulative bid quantities above the minimum generation for the units. For PDS schedule units this is cumulative PDS bid quantities.
PDS energy offer step quantity adjusted for minimum generation from OSM scheduled	*PDS_BidQS*	g, d, s	MW	Net bid step quantities that have minimum energy removed for OSM scheduled units in phases 3, 4, and 5. This is equal to *PDS_BidQ* for phases 1 and 2 and for PDS scheduled units.

[24] In the model, it is assumed that all FCAS units have 45° angles between a) their enablement minimum and low breakpoint, and b) their high breakpoint and enablement maximum.

PARAMETER	SYMBOL	SETS	UNITS	COMMENTS
units in phases 3, 4, and 5				
PDS energy offer step quantity used to calculate *PDS_BidQD*	*PDS_BidQC*	g, d, s	MW	Used to build PDS_BidQS and contains cumulative bid quantities. For OSM units this is cumulative bid quantities minue minimum generation for the units. For PDS schedule units this is cumulative PDS bid quantities.
OSM start-up cost	*OSMBid_SUCst*	g	$000	
OSM hourly running cost	*OSMBidRunCst*	g	$000/hr	
Hourly OSM running cost for synchronous condenser (and generating units running in synchronous condenser mode)	*Gen_SyncConHCst*	gc	$000/hr	
INITIAL CONDITIONS				
Initial resource status prior to start of run for the block	*InitGenStatus*			1 if on, 0 otherwise
First interval that resource can change status	*InitGenFI*			
Generation level prior to start of run for the block	*InitGenSG*	MW		
Phase 3 (OSM) synchronous mode flag by dispatch interval	*PDS_SolStatSynD*	*g,d*	(1 0n, 0 off)	This variable is set to 0 in phases 1, 2, and 3 and for phases 4 and 5 set to 1 if the unit is dispatch in synchronous mode during phase 3.
Phase 3 (OSM) synchronous mode flag but by capacity interval	*PDS_SolStatSynC*	*g,c*	(1 0n, 0 off)	Phase 3 synchronous mode flag by capacity interval like PDS_SolStatSynD above
If generation started in c, then number or capacity intervals it is active, 0 otherwise	*PDS_SolStatC*	g,c		

PARAMETER	SYMBOL	SETS	UNITS	COMMENTS
Generation offer status.	*PDS_SolStatD*	g,d		1 denotes using PDS bids and 0 denotes using OSM bids. For phase 1 and 2, these are all set to 1. For phase 3 (OSM), these will be set to 1 for units that are self-committed in the dispatch interval, and scheduled in phase 2 and for all IBR resources. For all other cases, will be set to 0. For phases 4 and 5, these will be set to 0 for all units scheduled as OSM units in phase 3 for the dispatch interval and 1 for all other cases.
PDS_SolStatD but by capacity interval	*PDS_SolStatS*	g,c		

Table 22 - Other parameters

PARAMETER	SYMBOL	SETS	UNITS	COMMENTS
Index of first dispatch interval *d* in commitment interval *c*	cStart	c		For example, if there are half-hourly dispatch intervals and two-hourly commitment intervals, then dispatch intervals 5, 6, 7 and 8 would correspond to commitment interval 2. cStart(c=2) = 5, and cEnd(c=2) = 8
Index of first dispatch interval *d* in commitment interval *c*	cEnd	c		See column above.
Local FCAS requirement flag	MnFCAS_SRF	sr		1 if a local FCAS requirement applies, zero otherwise.
Mapping from capacity intervals to dispatch intervals	cd	c, d		1 if dispatch interval d is part of capacity interval c.
Mapping of sub-regions to regions	sMapR	r, sr		1 if sub-region sr is part of region r.
Initial Storage volume for the block run	*InitStorVol*		MWh	
Required storage level at the end of the block run	*EndStorVol*		MWh	

PARAMETER	SYMBOL	SETS	UNITS	COMMENTS
Maximum imports for Victoria by input dispatch interval for phases 4 and 5	*PDS_VicMxImpz*	*dz*	MW	See below.
Maximum imports for Victoria by dispatch interval for phases 4 and 5	*PDS_VicMxImp*	*d*	MW	This variable is used in phases 1, 2, and 3 but set to a very large number so it is not binding. For Phases 4 and 5, this is set to the selected security configuration maximum import level.
Maximum wind and solar production for the sub-region	*WSCnst*	*sr,d*	MW	For all regions but SA and for all regions in phases 1, 2, and 3, this variable is set to a large number so the constraint is not binding. For SA in phases 4 and 5, this is set to the wind and solar production from phase 3.
Required capacity status for actual period of block run.	*tCapSP*	*g,c*	0 or 1	During the block running, the capacity status for the forecast period in the previous block is fixed in the current block. tCapSP contains these values.

Table 23 - Scalars

SCALAR	SYMBOL	UNITS	COMMENTS
First capacity interval for the phase and block	icMinCmp	Interval	
Last capacity interval for the phase and block	icMaxCmp	Interval	
First dispatch interval for the phase and block	idMinCmp	Interval	
Last dispatch interval for the phase and block	idMaxCmp	Interval	
Set to true if the PDS bids are scheduled with generator constraints	PDS_SchG	True/False	
Set to true if OSM bids are to be scheduled	OSM_SchG	True/False	
Set to 1 if either PDS_SchG or OSM_SchG specified	Either_SchG	1 or 0	
Number of dispatch intervals in capacity interval	iDiff	Interval	Must be constant for all capacity intervals.

A.3 Variables

Table 24 shows variables used in the OSM optimisation problem for each phase and optimisation stage (block). The variable column indicates the sets over which the variables are defined. The type column denotes the range of values the variable may take, e.g., binary (B), positive continuous (P) or positive integer (I).

An asterisk in the type column indicates that a variable was declared as being positive, but is restricted to taking binary values by constraints in the model. Such variables can be declared as being binary, but declaring them as positive may have performance benefits (as it reduces the number of binary variables while not affecting the optimal solution).

Table 24 - OSM optimisation variables

VARIABLE	DESCRIPTION	UNITS	TYPE
vCapOn(g, c)	Unit start-up status; 1 if capacity comes online at the start of the capacity interval, 0 otherwise		B*
vCapOff(g, c)	Unit shut-down status: 1 if capacity goes offline at the start of the capacity interval, 0 otherwise		B*
vCapC(g, c)	1 if the capacity status can change at the start of the capacity period and 0 otherwise		B*
vCapC0(g)	Initial capacity commitment status		B*
vGenS(g, d, s)	Scheduled PDS generation for generating unit, dispatch interval and bid step	MW	P
vGenSd0(g, d)	Scheduled PDS generation for unit, dispatch interval, and step 0 (OSM bid reset)	MW	P
vGenS0	Scheduled PDS generation for generating unit, dispatch interval=0, bid step, and stage, MW	MW	P
vTransF(tLnk, d)	Forward transmission from origin subregion to destination subregion	MW	P
vTransR(tLnk, d)	Reverse transmission from origin subregion to destination subregion	MW	P
vNetFCAS(sr, dir, f, d)	Net FCAS in the subregion	MW	P
vFCAS(g, dir, f, d, s)	Scheduled FCAS for generating unit, dispatch interval, bid step, FCAS direction and FCAS type	MW	P
vExcessDmd(sr, d)	Excess generation scheduled for region, dispatch interval, and stage	MW	P
vLessDmd(sr, d)	Shortfall in generation scheduled for regions, dispatch interval, and stage	MW	P
vExcessUpRmp(g, d)	Excess upward ramping for generator unit, dispatch interval, and stage	MW	P
vExcessDwnRmp(g, d)	Excess downward ramping for generator unit, dispatch interval, and stage	MW	P
vLessUnits(r, c)	Shortfall in required number of synchronous units	# units	I*
vLessInertia(r, c)	Shortfall in required inertia	MWs	P
vLessFCAS6(r, d)	Shortfall in regional 6sec FCAS	MW	P
vLessFCAS_SR(sr, dir, f, d)	Shortfall in required FCAS by type and subregion (only used if islanding condition creates demand for FCAS in sub-region).	MW	P
vLessFCAS_Glb(dir, f, d)	Shortfall in required global FCAS by type	MW	P

VARIABLE	DESCRIPTION	UNITS	TYPE
vLessCSUnits(csg, c)	Shortfall in minimum number of units for the configuration constraint	# units	I*
vSCnst_SA(so_SA, c)	Configuration security constraint option for SA		B
vSCnst_VIC(so_VIC, c)	Configuration security constraint option for VIC		B
vGenPDSDiscount(g, c)	Variable used to account for when an OSM is started, but this is only an extension of a self-commitment.	# starts	P
vCntPDSA(g,c)	Flag set to 1 if any PDS scheduled unit occurs in the current capacity interval and all following intervals to the next start for the unit. But flag set to 0 at all capacity intervals where a startup occurs		B*
vSAErr(c)	Error in assigning security options for SA for each capacity interval.		B*
vVicErr(c)	Error in assigning security options for VIC for each capacity interval.		B*
vStorVol(gx, d)	Storage volumes for batteries and pumped storage	MWh	P
vStorVol0(gx)	Initial storage volumes for batteries and pumped storage	MWh	P
vStorInj(gx, d)	Storage injection rate before injection losses for batteries and pumped hydro units, and dispatch interval	MW	P
vSynC(gc, c)	Measures the initiation of the conversion to SC mode. Set to 1 if unit converts to synchronous mode in the capacity interval.		B
vLSD(gx, d)	Shortfall in storage volumes for batteries and pumped storage in the dispatch interval potentially caused by block starting and ending storage constraints	MWh	
vHSD(gx, d)	Excess storage volumes for batteries and pumped storage in the dispatch interval potentially caused by block starting and ending storage constraints	MWh	
vTCost	Total cost of scheduled capacity and generation	$000	P
vGenG(g, d)	Total scheduled generation for the generating unit and dispatch interval.	MW	P
vGenG0(g)	Total scheduled generation for the generating unit at the start of the run.	MW	P
vGenSR(sr, d)	Total scheduled generation for each sub-region and dispatch interval.	MW	P
vFCAST(g, dir, f, d)	Scheduled FCAS for generating unit, dispatch interval, bid step, FCAS direction, FCAS type.	MW	P
vWSgen(sr, d)	Total IBR (wind and solar) generation for subregion and dispatch interval	MW	P
CCost(ccmp)	Total cost by cost component	$000	

VARIABLE	DESCRIPTION	UNITS	TYPE
CCostR(r, ccmp)	Total cost by region and cost-component	$000	
vCapS(g, c)	Set to 1 if unit is scheduled online during the capacity interval, 0 otherwise.		B
vCapS0(g)	Set to 1 if capacity is online during the capacity interval=0, 0 otherwise.		B
vSynS(gc, c)	Status for dual mode generating unts - 1 if capacity in synchronous mode and 0 otherwise		B
vSCnst_SA(so_SA, c)	Configuration security constraint option for SA		B
vSCnst_VIC(so_VIC, c)	Configuration security constraint option for VIC		B
vBatS(gb, c)	0 if battery is selected to provide inertia in the capacity interval, and 1 if not.		B

A.4 Objective Function

The objective of the OSM is to maximise the gains from trade. Qualitatively, this is simply the difference between the consumer value and the producer cost, and is the same objective that applies in the PDS, NEMDE, and any general market clearing problem.[25] However, the mathematical objective function in the OSM is not the same as for the PD, and there are some complexities in defining it. Furthermore, to allow the OSM to solve when there may be infeasibilities, the objective function includes several non-physical penalty costs on slack variables.

Hence, Table 25 presents the core objective function terms, while Table 26 describes additional penalty function terms (which in principle would not be needed if it could be guaranteed that a feasible solution will always exist).

Note that in the prototype, the objective function was expressed different so that terms were calculated by region and component, e.g., splitting out synchronous condenser costs from generation costs. This was useful for the purposes of results analysis, but is only a different way of calculating the same overall number.

Table 25 - Core objective function terms

DESCRIPTION	SETS	COMMENTS
Value of clearing demand (+)	d, r	All demand in the prototype is assumed non-scheduled, but price responsive scheduled demand could also be incorporated.
Cost of clearing PDS offers for energy (-)	g, d, s	
Cost of clearing PDS offers for FCAS (-)	g, d, f, f_d, s	
Cost of running cleared OSM units (-)	g, c,	
Cost of starting cleared OSM units (-)	g, c	

The penalty costs on slack variables in Table 26 are intended to allow the model to find a feasible solution when it is not possible to satisfy all required physical and power system constraints. These serve the same function as the constraint violation penalties (CVP) used in NEMDE.[26] The relative value of penalties should be set so as to guide the solver into which constraints should be prioritised (e.g., there may be a higher penalty on a higher quality ancillary service than a lower quality service). However, as constraints use different units (e.g., FCAS being measured in MW whereas inertia is measured in MW.sec) this can also mean that different penalty costs are used.

Table 26 - Objective function penalty terms ($000s)

SLACK VARIABLE	SETS
Excess up ramp (*vExcessUpRmp*) and down ramp (*vExcessDwnRmp*)	g, d
Unserved demand (*vLessDmd*) and overserved demand (*vExcessDmd*)	sr, d
Unmet local FCAS (vLessFCAS_*SR*)	sr, f, f_d, d
Unmet global FCAS (*vLessFCAS_Glb*)	f, f_d, d
Insufficient synchronous units (*vLessUnits*)	r, c

[25] In a market clearing problem, the gains from trade is the difference between the value of the produced goods (e.g., energy) to consumers, less the cost of producing those goods from producers.
[26] Refer to AEMO – Schedule of Constraint Violation Penalty Factors, 2017.

SLACK VARIABLE	SETS
Insufficient inertia (*vLessInertia*)	r, c
Units missing from configuration constraint station grouping (*vLessCSUnits*)	c, csg
Security configuration not met in SA (*vSAErr*) and Victoria (*vVicErr*)	c

A.5 Constraints

The following sections set out the constraints used in the phase 3 GAMS model. Each table lists the label used in the GAMS model, the indices for which the constraint is defined over, a short description, and the constraints itself.

Where there are multiple tables with the same constraint label, and the reference uses the same number, this is because one GAMS constraint covers multiple constraint forms by using conditional terms. For example, constraint 1A and 1B (*cCapS*) are defined in the same constraint, with 1B being an initial interval version of 1A (so it refers to an initial parameter, rather than a variable in the previous interval).

Operating Mode Status Constraints

The following constraints control operating mode transitions – e.g., from online to offline, or from synchronous condenser mode to synchronous generation mode.

Ref: 1A	**Label:** cCapS	**Indices:** $g \in G; \quad c \in C, c > 1$
Description	Online capacity equation of state; capacity status is equal to capacity status in previous commitment interval, adjusted for any start-ups or shut-downs.	
Equation	$vCapS(g,c) = vCapS(g,c-1) + vCapOn(g,c) - vCapOff(g,c)$	

Ref: 1B	**Label:** cCapS	**Indices:** $g \in G; \quad c = 1$
Description	Online capacity equation of state for the first interval. As for Eq. 1A (cCapS), but uses the initial capacity status in place of the capacity status of the previous interval.	
Equation	$vCapS(g,c) = vCapS0(g) + vCapOn(g,c) - vCapOff(g,c)$	

Ref: 2A	**Label:** cCapC	**Indices:** $g \in G; \quad c \in C, c > 1$
Description	Controls whether the capacity status can change (CapC = 1) or not (CapC=0). Note, the last two terms are only included if $c_{ON} > 0$ and $c_{OFF} > 0$ respectively.	
Equation	$vCapC(g,c) = vCapC(g,c-1) + vCapOn(g,c) - vCapOff(g,c) + vCapOn(g,c_{ON}) + vCapOff(g,c_{OFF})$ Where $c_{ON} = c - MnOn(g,c)$ and $c_{OFF} = c - MnOff(g,c)$	

Ref: 2B	**Label:** cCapC	**Indices:** $g \in G; \quad c = 1$
Description	As for 2A, but uses the initial capacity status in place of the capacity status of the previous interval. Note, the last two terms are only included if $c_{ON} > 0$ and $c_{OFF} > 0$ respectively.	
Equation	$vCapC(g,c) = vCapC0(g) + vCapOn(g,c) - vCapOff(g,c) + vCapOn(g,c-c_{ON}) + vCapOff(g,c-c_{OFF})$	

	Where $c_{ON} = c - MnOn(g,c)$ and $c_{OFF} = c - MnOff(g,c)$

Ref: 3	**Label:** cSynS	**Indices:** $gc \in G_C; \quad c \in C$
Description	Requires that a dual mode generating unit is not simultaneously in synchronous generation mode (vCapS = 1) and synchronous condenser mode (vSynS=1).	
Equation	$vSynS(gc,c) \leq 1 - vCapS(gc,1)$	

Ref: 4	**Label: cSynSD**	**Indices:** $gc \in G_C\ st.Gen\_SynCvnDly = 1; \quad c \in C, c > 1$
Description	Requires that dual mode generator with a switching delay is not in synchronous generation mode (CapS = 1) the interval following an interval in which it was in synchronous condenser mode (CapS=1). Note: This assumes the switching delay is equal to one period.	
Equation	$vSynS(gc,c) \leq 1 - vCapS(gc,c-1)$	

Ref: 5	**Label:** cVSynSOff	**Indices:** $gc \in G_C; \quad c \in C$
Description	Requires that unit cannot be schedule in synchronous mode during the minimum off time after the unit is shut down	
Equation	$vSynS(gc,c) \leq 1 - vCapC(gc,c)$	

Ref: 6	**Label:** cSynC	**Indices:** $gc \in G_C; \quad c \in C, c > 1$
Description	Synchronous condenser mode equation of state.	
Equation	$vSynC(gc,c) \geq vSynS(gc,c) - vSynS(gc,c-1)$	

Extended and Avoided Starts Constraint

Ref: 7	**Label:** cCapDisc	**Indices:** $g \in G; \quad c \in C$
Description	Restricts cCapDisc to the value of 1 (which in turn allows unit g's start-up cost to be subtracted in the objective function) if the unit is OSM-committed for the period in which it is self-committed.	
Equation	$vGenPDSDiscount(g,c) + vCntPDSA(g,c) \leq vCntPDSA(g,c+1) + vCapS(g,c) * PDS\_SolStatS(g,c)$	

Ref: 8	**Label:** cCapDiscX	**Indices:** $g \in G; \quad c \in C$
Description	Restricts the capacity start-up discount to the sum of the OSM unit capacity start-ups over the day or block periods"	
Equation	$vGenPDSDiscount(g,c) + vCntPDSA(g,c) \leq vCntPDSA(g,c+1) + vCapS(g,c) * (PDS\_SolStatS(g,c)$	

Ref: 9	Label: cCapDiscX	Indices: $g \in G$
Description	Restricts the capacity start-up discount to the sum of the OSM unit capacity start-ups over the day or block periods (Only for phase 3)	
Equation	$\sum_{c \in C} vGenPDSDiscount(g,c) \leq \sum_{c \in C} vCapOn(g,c) \quad \text{where } PDS_{SolStatS(g,c)} = 0$	

Intermediate generation output and total sub-region generation variables

Ref: 10	Label: cTGen	Indices: $sr \in S_R; \quad d \in D$
Description	Calculates total generator output as the sum of the cleared PD offers, minimum generation if OSM commited, and for phases 4 and 5 - the generation from the 1st bid step (step 0) representing the OSM bid in the PDS	
Equation	$vGenG(g,d) = \sum_s vGenS(g,d,s) + \sum_{\substack{c \\ \text{st. } cd(c,d)}} (1 - PDS\_SolStatS(g,c)) \times vCapS(g,c) \times GenMnGen(g) +$ $vGenSd0(g,d)$	

Ref: 11	Label: cTGenSR	Indices: $g \in G; \quad d \in D$
Description	Calculates total generator output for each subregion	
Equation	$vTGenSR(sr,d) = \sum_{g \in G_{RS}} vGenG(g,d)$	

Minimum and maximum output

Ref: 12	Label: cGenMx	Indices: $g \in G; \quad d \in D$
Description	Requires generation from unit *g* to be less than its capacity if committed, or zero otherwise.	
Equation	$vGenG(g,d) \leq GenMxGen(g) \times \sum_{\substack{c \\ \text{st. } cd(c,d)}} vGenG(g,d)$	

Ref: 13	Label: cGenMn	Indices: $g \in G; \quad d \in D$
Description	Requires generation from unit *g* to be greater than its minimum generation if committed, or zero otherwise.	
Equation	$vGenG(g,d) \geq GenMnGen(g) \times \sum_{\substack{c \\ \text{st. } cd(c,d)}} vGenG(g,d)$	

Sub-region energy balance

Ref: 14	Label: cDmd	Indices: $sr \in R_s; \quad d \in D$
Description	Nodal energy balance, consisting of i) sub-regional demand ii) transmission flows for lines flowing away from the sub-region iii) transmission flows for lines flowing towards the sub-region iv) storage injections (to the unit from the grid) and v) slack variables.	
Equation	$rDmd(sr,d) = vGenSR(sr,d) - \sum_{\substack{t\ st.\\ t\in T_{f,SR}}} [vTransF(t,d) - vTransR(t,d) \times (1 - Trans\_Loss(t))] - \sum_{\substack{t\ st.\\ t\in T_{t,SR}}} [vTransR(t,d) - vTransT(t,d) \times (1 - Trans\_Loss(t))] + \sum_{\substack{gx\ st.\\ gx\in G_{SR}}} vStorInj(gx,d) + vLessDmd(sr,d) - vExcessDmd(sr,d)$	

Ramp rate constraints

Ref: 15A	Label: cRmpU	Indices: $g \in G; \quad d \in D, d > 1$
Description	Ramp up constraint. The last term ensures that when a unit is turned on, its ramp rate does not limit it getting above minimum generation (although GenMxGen has been used in the accompanying GAMS code).	
Equation	$vGenG(g,d) - vGenG(g,d-1) - vExcessUpRmp(g,d) \leq GenRmpUp(g) + \sum_{c.st.cStart(c)=d} vCapOn(g,c) \times GenMxGen(g)$	

Ref: 15B	Label: cRmpU	Indices: $g \in G; \quad d \in D, d = 1$
Description	Ramp up constraint for initial interval	
Equation	$vGenG(g,d) - vGenG0(g) - vExcessUpRmp(g,d) \leq GenRmpUp(g) + \sum_{c.st.cStart(c)=d} vCapOn(g,c) \times GenMxGen(g)$	

Ref: 16A	Label: cRmpD	Indices: $g \in G; \quad d \in D, d > 1$
Description	Ramp dow constraint. The last term ensures that when a unit is turned on, its ramp rate does not limit it getting to zero (although GenMxGen has been used in the accompanying GAMS code, it might would be possible to have a lower value if desired).	
Equation	$vGenG(g,d-1) - vGenG(g,d) - vExcessDwnRmp(g,d) \leq GenRmpUp(g) + \sum_{c.st.cStart(c)=d} vCapOff(g,c) \times GenMxGen(g)$	

Ref: 16B	Label: cRmpD	Indices: $g \in G; \quad d \in D, d = 1$

Description	Ramp down constraint for initial interval
Equation	$vGenG0(g) - vGenG(g,d) - vExcessDwnRmp(g,d) \leq GenRmpDwn(g) + \sum_{c.st.cStart(c)=d} vCapOff(g,c) \times GenMxGen(g)$

FCAS constraints

Ref: 17	**Label:** cFCAS_Tot	**Indices:** $g \in G, fd \in F_d, f \in F, d \in D$
Description	Calculates total cleared FCAS quantity (service *f*, direction *d*) for unit *g*, dispatch interval *d* as the sum of the cleared quantity in each of the ten FCAS offer bands.	
Equation	$vFCAST(g,fd,f,d) = \sum_{s \in S} vFCAS(g,fd,f,d,s)$	

Ref: 18	**Label:** cFCAS_Avl	**Indices:** $g \in G, fd \in F_d, f \in F, d \in D$
Description	Limits total cleared FCAS quantity (service *f*, direction *d*) for unit *g*, dispatch interval *d* to the FCAS capability if committed, zero otherwise.	
Equation	$vFCAST(g,fd,f,d) \leq FCAS\_Cap(g,fd,f) \times vCapS(g,c=cd(c,d))$	

Ref: 19	**Label:** cFCAS_LimitR	**Indices:** $g \in G, f \in [6\text{s}, 60s, 5min], d \in D$
Description	Requires that there is sufficient headroom above energy dispatch to provide the combined quantity of each raise contingency FCAS and raise regulation	
Equation	$vGenG(g,d) + vFCAST(g,fd=\text{Raise, f, d}) + vFCAST(g,fd=\text{Raise, "Reg", d}) \leq vCapS(g,c=cd(c,d)) \times GenMxGen(g)$	

Ref: 20	**Label:** cFCAS_LimitL	**Indices:** $g \in G, f \in [6\text{s}, 60s, 5min], d \in D$
Description	Requires that there is sufficient footroom below energy dispatch to provide the combined quantity of each lower contingency FCAS and lower regulation	
Equation	$vGenG(g,d) - vFCAST(g,fd=\text{Lower, f, d}) \text{-} vFCAST(g,fd=\text{Lower, "Reg", d}) \geq vCapS(g,c=cd(c,d)) \times GenMnGen(g)$	

Ref: 21	**Label:** cFCAS_Net	**Indices:** $sr \in R_s, f \in F, fd \in F_d, d \in D$
Description	Calculates total quantity of cleared FCAS (service *f*, direction *fd*) in sub-region *sr*.	
Equation	$vNetFCAS(sr,fd,f,d) = \sum_{g \in G_{SR}} vFCAST(g,fd,f,d)$	

Ref: 22	**Label:** cFCAS_Min_SR	**Indices:** $sr \in R_s\ st.MnFCAS_{SRF} = 1, fd \in F_d, f \in F, d \in D$

Description	For regions with a local FCAS requirement, requires that cleared FCAS in that sub-region meets that requirement.
Equation	$vNetFCAS(sr, fd, f, d) + vLessFCAS\_SR_{(sr,fd,f,d)} \leq MnFCAS\_SR_{(sr,d,fd,f)}$

Ref: 23	**Label:** cFCAS_Min_Glb	**Indices:** $fd \in F_d, f \in F, d \in D$
Description	Requires that the sum of cleared FCAS in each sub-region without a local requirement meets the global FCAS requirement.	
Equation	$\sum_{\substack{sr\ st.\\ MnFCAS\_SRF(sr,d) \leq 1}} vNetFCAS(sr, fd, f, d) + vLessFCAS\_Glb_{(fd,f,d)} \leq MnFCAS\_Glb_{(d,fd,f)}$	

Inertia and minimum synchronous unit constraints

Ref: 24	**Label:** cRInertia	**Indices:** $r \in R, c \in C$
Description	Requires that the inertia requirement (which may be zero) is met in each region. The right hand side is set to less than 0 for phases 1, 2, 4, and 5 so that the constraint is not binding.	
Equation	$\sum_{g \in G_R} [vCapS(g, c) \times GenInertia(g)] + \sum_{g \in G_C \cap G_R} [vSynS(g, c) \times GenInertia(g) \times GenInertia(g) \times Gen_{SynCvnIS(g)}] - \sum_{g \in Gb \cap G_R} [vCapS(g, c) \times GenInertia(g) \times vBatS(g, c)] + vLessInertia(r, c) \geq rInertia(r)$	

Ref: 25	**Label:** cRUnits	**Indices:** $r \in R, c \in C$
Description	Requires that minimum number of large synchronous units (which may be 0) is met in each region. The right hand side is set to less than 0 for phases 1, 2, 4, and 5 so that the constraint is not binding.	
Equation	$\sum_{g \in G_R \cap G_S} [vCapS(g, c) \times Genlrgsyn(g)] + vLessUnits(r, c) \geq rSUnits(r)$	

Ref: 26	**Label:** cRUnits_SynCon	**Indices:** $r \in R, c \in C$
Description	Requires a minimum number of synchronous condenser only units scheduled for the region and capacity interval	
Equation	$\sum_{g \in G_R \cap G_S} [vCapS(g, c) \times Gensynflg(g)] + vLessUnits(r, c) \geq rSUnits(r)$	

Configuration constraints

Ref: 27	**Label:** cSCnst_Sel_SA	**Indices:** $c \in C$
Description	Requires that one, and only one, valid configuration be selected (*vSCnst_SA(so_SA, c) = 1*) for South Australia. RHS is set to 0 in all phases except phase 3.	
Equation	$\sum_{so_{SA} \in SO_{SA}} vSCnst\_SA(so\_SA, c) + vSAErr(c) = 1$	

Ref: 28	**Label:** cSCnst_Sel_VIC	**Indices:** $c \in C$
Description	Requires that one, and only one, valid configuration be selected (*vSCnst_VIC(so_VIC, c) = 1*) for Victoria. RHS is set to 0 in all phases except phase 3.	
Equation	$\sum_{so_{VIC} \in SO_{VIC}} vSCnst\_VIC(so\_VIC, c) + vVICErr(c) = 1$	

Ref: 29	**Label:** cSCnst_MU_SA	**Indices:** $csg \in CSG, c \in C$
Description	Ensures that at least *sCnst_MinUnits_SA* from constraint station group *csg* are committed when configuration *so_SA* is active. While the GAMS code includes the slack variable vLessCSUnits, this could be removed as there is a slack variable in constraint *cSCnst_Sel_SA*. Also, in phases 1, 2, 4 and 5, since vSCnst_SA(so_SA, c) is 0 for all configurations, this constraint will not be binding.	
Equation	$\sum_{g} gMapcsg(g, csg) \times vCapS(g, c) + vLessCSUnits(csg, c) \geq \sum_{so_{SA}} [vSCnst\_SA\,(so\_SA, c) \times SCnst\_MinUnits\_SA(so\_SA, csg)]$	

Ref: 30	**Label:** cSCnst_MU_VIC	**Indices:** $csg \in CSG, c \in C$
Description	Ensures that at least *sCnst_MinUnits_VIC* from constraint station group *csg* are committed when configuration *so_VIC* is active. While the GAMS code includes the slack variable vLessCSUnits, this could be removed as there is a slack variable in constraint *cSCnst_Sel_VIC*. Also, in phases 1, 2, 4 and 5, since vSCnst_VIC(so_VIC, c) is 0 for all configurations, this constraint will not be binding.	
Equation	$\sum_{g} gMapcsg(g, csg) \times vCapS(g, c) + vLessCSUnits(csg, c) \geq \sum_{so_{SA}} [vSCnst\_VIC(so\_VIC, c) \times SCnst\_MinUnits\_VIC(so\_VIC, csg)]$	

Ref: 31	**Label:** cSCnst_MxWS	**Indices:** $d \in D$

Description	Requires that dispatched IBR in SA be less than the allowable IBR associated with the active configuration.
Equation	$\sum_{g \in G_R \cap G_{IBR}} vGenG(g,d) \leq \sum_{c\ st.cd(c,d)=1} \sum_{so_{SA}} [vSCnst_{SA}(so_{SA}, c) \times SCnst\_MxWS(so\_SA)]$

Ref: 32	**Label:** cSCnst_MxImp	**Indices:** $d \in D$
Description	Requires that imports to Victoria be less than the allowable IBR associated with the active configuration. Note that almost all Victorian configurations have no import limit, and this constraint is thought not to have a material effect on most cases. It is included here for completeness only.	
Equation	$\sum_{t\ st.t \in T_{t,VIC}} vTransF(t,d) + \sum_{t\ st.t \in T_{f,VIC}} [vTransR(t,d) \times (1 - Trans_{Loss(t)})] \leq \sum_{c\ st.cd(c,d)=1} \sum_{so\_SA} [vSCnst\_VIC(so\_VIC, c) \times SCnst\_MnImp(so\_VIC)]$	

Ref: 33	**Label:** cSCnst_MxImpA	**Indices:** $d \in D$
Description	Requires that imports to Victoria be less than the allowable IBR associated with the active configuration selected in phase 3. mxImp(d) is set to a large number in phases 1, 2, and 3 so the constraint is not binding and set to the maximum imports from the configuration selected in phase 3 for phases 4 and 5	
Equation	$\sum_{t\ st.t \in T_{t,VIC}} vTransF(t,d) + \sum_{t\ st.t \in T_{f,VIC}} [vTransR(t,d) \times (1 - Trans_{Loss(t)})] \leq mxImp(d)$	

Intermittent resource availability constraints

Ref: 34	**Label:** cWSSm	**Indices:** $sr \in R_S, d \in D$
Description	Calculates the total variable (wind and solar) generation in a sub-region.	
Equation	$vWSgen(sr,d) = \sum_{g \in G_{IBR} \cap G_{SR}} vGenG(g,d)$	

Ref: 35	**Label:** cWSCnst	**Indices:** $sr \in R_S, d \in D$
Description	Limits the combined wind and solar in a region to be less than the scheduled generation in phase 3. For phases 1, 2, and 3.WSCnst(sr, d) is set to a large number so the constraint is not binding. For phases 4 and 5, WSCnst(sr, d) is set to the scheduled wind and solar generation from phase 3.	
Equation	$vWSgen(sr,d) \leq WSCnst(sr,d)$	

Storage constraints

Note that these equations were included in the prototype, but would not be required in the OSM if the storage participant is expected to manage their own storage levels.

Ref: 36A	**Label:** cStorVols	**Indices:** $gx \in G_X;\quad d \in D, d > 1$
Description	Tracks stored energy for batteries and other storage units. Note that storage injection is before losses and storage withdrawal (vGenG) is after losses.	
Equation	$vStorVol(gx,d) = vStorVol(gx,d-1) + vStorInj(gx,d) \times (1 - \frac{GenInjLoss(gx)}{2}) + vGenG(gx,d) \times (1 - \frac{GenWthLoss(gx)}{2})$	

Ref: 36B	**Label:** cStorVols	**Indices:** $gx \in G_X;\quad d = 1$
Description	Tracks stored energy for batteries and other storage units – initial interval. Note that storage injection is before losses and storage withdrawal (vGenG) is after losses.	
Equation	$vStorVol(gx,d) = vStorVol0(gx) + vStorInj(gx,d) \times (1 - \frac{GenInjLoss(gx)}{2}) + vGenG(gx,d) \times (1 - \frac{GenWthLoss(gx)}{2})$	

Battery Inertia

Ref: 37	**Label:** cBatInertiaL	**Indices:** $gb \in G_B, d \in D$
Description	Restricts battery injection to 90% of injection capacity if battery used for inertia.	
Equation	$vStorInj(gb,d) \leq GenMxInjRte(g) * BatIneritaLim + \sum_{c\ st.cd(c,d)=1} vBatS(gc,c) \times 999999.9$	

Ref: 38	**Label:** cBatInertiaH	**Indices:** sr, d
Description	Restricts battery withdrawal to 90% of withdrawal capacity if battery used for inertia	
Equation	$vGenG(gb,d) \leq GenMxGen(g) * BatIneritaLim + \sum_{c\ st.cd(c,d)=1} vBatS(gc,c) \times 999999.9$	

A.6 Notes

Because the prototype is a toy model, it includes a number of features that would be modified or unneeded if implemented in the actual NEM scheduling process. This section provides some commentary on these and other aspects:

- Constraints 36 (cStorVols) was used in the OSM prototype (and other phases) to track storage volumes and to ensure that storage devices are required to recharge any energy that is discharged (including losses). These constraints avoid infeasible behaviour, but in the NEM dispatch, the participant is required to adjust its offers so that they are consistent with their unit's constraints (although they can also apply daily energy limits). Given the OSM is an inter-temporal problem, this warrants consideration. Currently, it is thought that a daily energy limit would be sufficient for a production version OSM, but this is to be given further consideration in the course of this work.

- Some basic constraints are included as upper bounds that are applied to variables. For example:
    - Transmission flow upper and lower bounds are set to the transmission capacity.
    - Storage volumes upper bounds are set to the storage energy capacity. The upper bound on the storage charge rate is also set to the maximum charge rate (after adjusting for losses).
    - The upper bound on the generation from wind and solar units is set to their forecast resource availability (i.e., the UIGF).