

# CONSORT

## Bruny Island Battery Trial

### Project Final Report

# Network-Aware Coordination (NAC)

Paul Scott (contact author), Dan Gordon, Sylvie Thiébaux  
The Australian National University

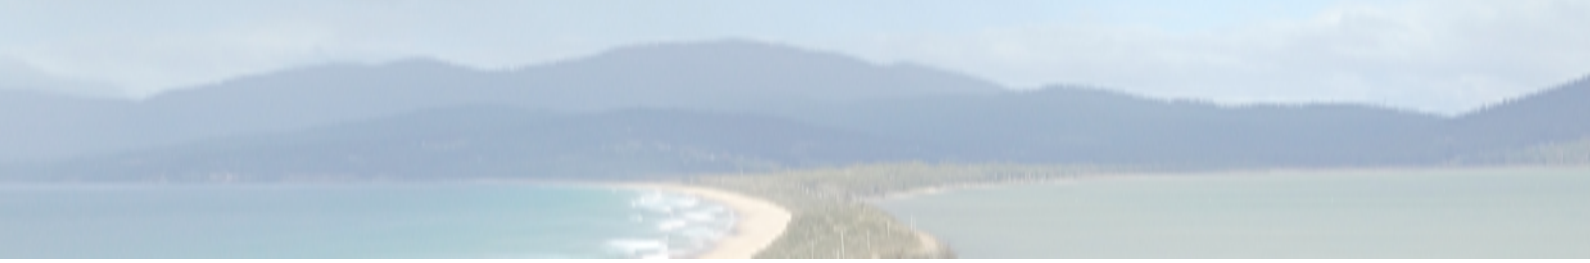
[Paul.Scott@anu.edu.au](mailto:Paul.Scott@anu.edu.au)

April 2019

<http://brunybatterytrial.org>

CONSORT: “Consumer Energy Systems Providing Cost-Effective Grid Support” is a collaboration between The Australian National University, The University of Sydney, University of Tasmania, Reposit Power and TasNetworks. The Australian Government, through the Australian Renewable Energy Agency, is providing \$2.9m towards the \$8m trial under its Research and Development Program.







## Executive Summary

Network-Aware Coordination (NAC) is the key technological innovation at the heart of the CONSORT project. Over the three years of the project, NAC has moved from a stage of basic research, to a design and robust implementation that runs in live network operations.

NAC consists of algorithms, techniques, and software for automatically coordinating distributed energy resources (DER) in a way that both respects network constraints, and minimises the total cost to the network provider and the DER owners. In effect, it achieves *optimal power flow* for the distribution system. The NAC preserves participants' privacy and agency, and implements a distributed algorithm that enables it to scale to larger problems.

The CONSORT project has demonstrated the benefits of NAC on the Bruny Island network, by using NAC-orchestrated DER, in this case customer-owned solar and batteries, to reduce the island's reliance on expensive diesel generation, which was previously TasNetworks' only existing means of managing network constraints. The project has reduced the diesel requirements by around a third for the trial period. In addition to this, through detailed simulations that complement the live trials, we have demonstrated that this reduction is not just an artifact of having more DER installed on the island, but that the NAC brings significantly better outcomes over how the DNSP would manually dispatch the DER.

From the beginning, the NAC technology has been designed to not just solve the Bruny Island problem, but to also expand to a larger range of network conditions, to different types of DER including electric vehicles, and to enable load flexibility to the wider wholesale markets. It provides the means for distributed renewable generation and other DER to be safely, efficiently and flexibly integrated into our distribution networks, improving the outcomes for network operators, DER owners and the wider NEM. We propose future NAC development and trials to focus on delivering and demonstrating this full range of capabilities, making the smart grid an everyday reality.



# Table of Contents

<b>Executive Summary</b>	<b>3</b>
<b>Table of Contents</b>	<b>4</b>
<b>1 Introduction</b>	<b>6</b>
<b>2 The Bruny Island Feeder Problem</b>	<b>6</b>
2.1 Beyond Bruny Island	7
<b>3 The NAC Approach</b>	<b>8</b>
3.1 Design	8
3.1.1 NAC prices	8
3.1.2 Rolling horizon optimisation	9
3.1.3 The ADMM algorithm	9
3.2 Implementation	10
3.2.1 Detailed architecture	10
3.2.2 Reposit boxes (remotes)	11
Proxy remotes	11
3.2.3 NAC workers	12
3.2.4 The NAC dealer	12
3.2.5 REDIS in-memory database	13
3.2.6 e-JSON data format and the network model	13
Network phasing	13
Phasing of participant houses	14
Open-delta voltage regulator	14
3.2.7 Load prediction	14
<b>4 Simulations</b>	<b>15</b>
4.1 CONSORT simulation architecture	15





<b>5 Trials</b>	<b>16</b>
5.1 Goals from NAC perspective	16
5.2 Relevant trial periods	17
5.3 Supplementary simulations	17
<b>6 Results</b>	<b>18</b>
6.1 Summary of NAC trial periods	18
6.2 Details of trial periods	19
6.2.1 Easter 2018 (TP180329)	19
6.2.2 April school holidays (TP180413)	20
6.2.3 Queen's Birthday long weekend (TP180608)	20
6.2.4 July school holidays (TP180713)	22
6.2.5 Summer holidays 2018-19 (TP20181219)	23
6.2.6 Australia Day (TP190125)	24
6.2.7 Voltage control tests (TP190130)	24
6.2.8 Shapley value payments (TP190208)	24
6.3 Simulations and further analysis	25
6.4 Voltage control	28
6.5 Performance and computational burden of NAC	30
<b>7 Conclusion</b>	<b>31</b>
7.1 Recommendations going forward	32
7.1.1 Voltage control	32
7.1.4 Tools for discovering network state and configuration	32
7.1.2 Scaling	32
7.1.3 Market integration	33
<b>References</b>	<b>33</b>



# 1 Introduction

NAC, or Network Aware Coordination (NAC), is a means to coordinate distributed energy resources (DER) in a way that respects network constraints and minimises the total cost to both the network provider and the DER owners (participating households). NAC is the key innovation at the heart of the CONSORT project. It has the following desirable properties:

- In comparison with competing approaches to coordination, it actually solves the problem (ie it does respect network constraints and it does so at minimised cost).
- It is a distributed algorithm. As more NAC participants are added, the calculation can be scaled up by simply adding more computational resources, and leveraging the resources of the participants.
- It preserves the privacy of participants: they need only provide very limited information about their intended level of battery discharge over the next 24 hours.
- It provides the means for participants to have agency. NAC is not concerned with the details of the calculation by which each individual household energy management system (EMS) plans its battery discharge. This means that there is a clean interface between the NAC and the EMSs; each of these may be developed independently, and to participate in NAC, an EMS need only to be able to exchange some simple information, and to conform to some assumptions about the way the battery discharge is planned.

Over three years of the CONSORT project, we have successfully implemented and tested NAC for up to 32<sup>1</sup> participating households on Bruny Island. Each of these households is equipped with a PV / battery system that are controlled by a [Reposit Box](#) [1]. At the time of writing this report, we are able to run NAC on a routine basis, for periods of at least several weeks at a time, and have demonstrated a reduction in diesel usage during peak periods of around 34% [2].

## 2 The Bruny Island Feeder Problem

Bruny Island is a small island off Tasmania, with over 1000 dwellings. Two undersea cables supply electricity, with the southern cable, see Fig. 1, supplying over 800 customers via a long rural feeder. This cable reaches its thermal limit during some peak load events on the island, and voltages at the southern end of the island can also drop to their lower limits. TasNetworks currently manages this problem by dispatching a mobile diesel generator during exceptional network peaks [3].

---

<sup>1</sup> The final number of participants in the trial was 34. Due to hardware issues and the late commissioning of one of the participants because of factors beyond our control, the NAC trials reported on here contained either 31 or 32 participants, depending on the particular trial period.



The southern feeder has around 337 buses (once reduced). It has a peak load of approximately 1.3 MVA, and may experience significant phase unbalance, up to around 15% difference in current between phases.

Bruny Island was felt to be a good location for the CONSORT trial, due to there being a well-defined problem with the network (congestion in the undersea cable), and also because it being an island with a well-defined community and network lent itself to both the social research and technical aspects of the project.

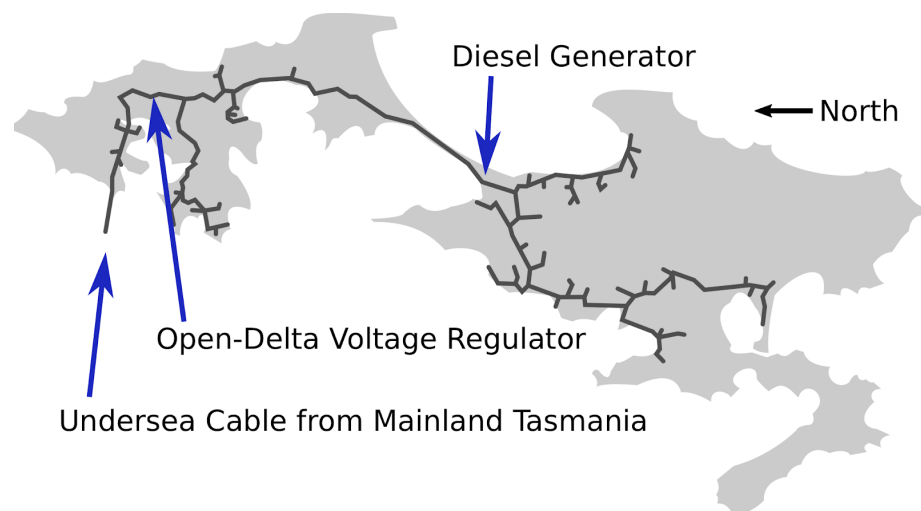


Fig. 1. The Bruny Island 11 kV network model used in CONSORT. Source: P. Scott *et al.* (2019) [2].

## 2.1 Beyond Bruny Island

While the cable constraint has provided a convenient focus for the CONSORT project, the NAC has been built in such a way to handle myriad other network constraints. Toward the end of the project, we investigated the use of NAC to control voltages.

Similarly, NAC was designed to be able to interface with multiple DER providers. As such, it can be adapted to work with a range of sources of DER - for example, different aggregators, VPPs, embedded generators, electric vehicles, etc. That said, there are currently restrictions on the way that providers can respond to prices: in a certain technical sense, they must act “rationally”, and the underlying mathematical problem they solve locally must be “well behaved”.

The NAC is therefore applicable over a wide range of feeders or networks of various sizes, experiencing a range of existing or emerging issues associated with DER deployment.

## 3 The NAC Approach

### 3.1 Design

The figure below shows the high-level design of NAC, as implemented in the CONSORT project. At the heart of the NAC system is the “NAC dealer”, whose primary responsibility is to oversee the convergence of the algorithm via a series of mediated negotiations involving the “remotes” (participating households) and the “workers” (network solvers).

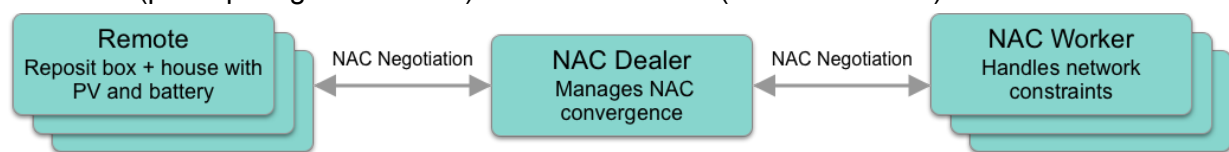


Fig. 2. Fundamental architecture of the NAC.

#### 3.1.1 NAC prices

NAC negotiations are based on a set of “NAC prices” - the price per kWh and kVarh which each house is offered for exporting real and reactive power to the network, or the amount that they are charged if they consume power.

There is a different real and reactive NAC price for each participating house, which account for network losses and constraints to each network node. These are often known as *locational marginal prices*, which are analogous to the wholesale market marginal spot prices. Similarly to the wholesale market, the NAC price can conceivably be negative, meaning that the houses are paid if they consume power and charged if they produce it. This could happen, for example, if the network voltages were too high and the NAC workers were trying to reduce voltages by increasing electricity consumption.

The NAC prices are not simply ad-hoc market prices. They are mathematically principled in the sense that, when the NAC negotiations have successfully concluded, the batteries’ response to the prices will correspond to the lowest total cost for the forward horizon.

It is important to note that in the CONSORT project, the prices which were negotiated do not correspond to actual payments and charges to participants. For a variety of reasons, the participants were paid for network support according to different rules, see Reward Structures Final Report [7]. They were paid \$1/kWh for network support, either in advance (energy reserve), or after the fact (energy use).

We also remark that in the CONSORT project, we have been concerned with the real power prices only and have to date considered the treatment of reactive power to be beyond the



scope of the project, due in part to technical work that needs to be done regarding inverter capabilities.

### 3.1.2 Rolling horizon optimisation

To achieve a satisfactory result, NAC must do forward planning, based on the best information available at the current time. It must continually update its plan, based on changing circumstances, and must always act according to the most recent plan. All of this is handled using the technique of “receding horizon optimisation”.

Every 5 minutes the NAC is run to find the cheapest plan covering approximately 24 hours. 48 timesteps are defined, as illustrated in Fig. 3 below. The first timestep starts approximately 5 minutes in the future, aligned to a five-minute boundary (e.g. 10:15). Additional timesteps occur at half-hour boundaries. Each such sequence of 48 timesteps is a “horizon”. Fig. 3 shows seven consecutive horizons.

The NAC has approximately five minutes to find a solution for each horizon. Upon doing so, it locks in the current solution for action by the participating households. It then waits until the next horizon, on the next five minute boundary, and starts the process again. Because each horizon overwrites the preceding one, under normal circumstances the batteries end up only ever acting on the first five minutes in each horizon.

However, if something goes wrong, for example if there is a communications outage or the NAC is not able to converge in the allocated five minutes, then some or all of the batteries may end up acting on older horizons.

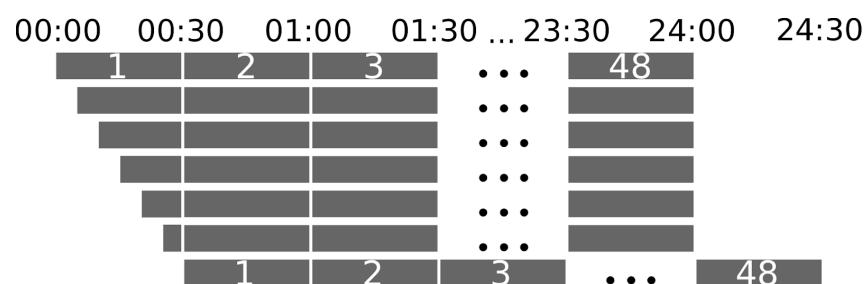


Fig. 3. Seven consecutive receding horizons at five minute intervals. Source: P. Scott *et al.* (2019) [2]

### 3.1.3 The ADMM algorithm

Mathematically, NAC solves an optimisation problem that ensures all constraints in the system are met at least total cost to the network and all participants. This problem is, loosely speaking, a version of the well known *multi period optimal power flow (OPF)* problem. It is a hard problem, due to the “non-convex” nature of several of the constraints.



The ADMM algorithm [4][2], as implemented in NAC, uses a “divide and conquer” approach to solve this problem by splitting it into a series of sub-problems.

Each remote (household) gets a single subproblem per horizon: to minimise the total cost (tariffs + NAC prices) by scheduling the battery, subject to constraints such as limits on the power and capacity limits for the battery.

Additionally, for each timestep in the horizon there is a worker subproblem: to minimise the total cost to the network by scheduling batteries and diesel dispatch, subject to network constraints such as power balance equations at nodes, thermal limits on lines, voltage upper and lower bounds, and physical power flow laws.

At each iteration of the ADMM algorithm, the subproblems are solved by the remotes and workers. The solutions and prices are then compared, and if the total discrepancy in the prices or dispatch levels is greater than a given tolerance, another iteration of the algorithm will be triggered, with suitably adjusted prices. Otherwise, the algorithm is deemed to have successfully converged.

## 3.2 Implementation

### 3.2.1 Detailed architecture

Fig. 4, below, gives a more detailed view of the NAC architecture as implemented in the CONSORT project.

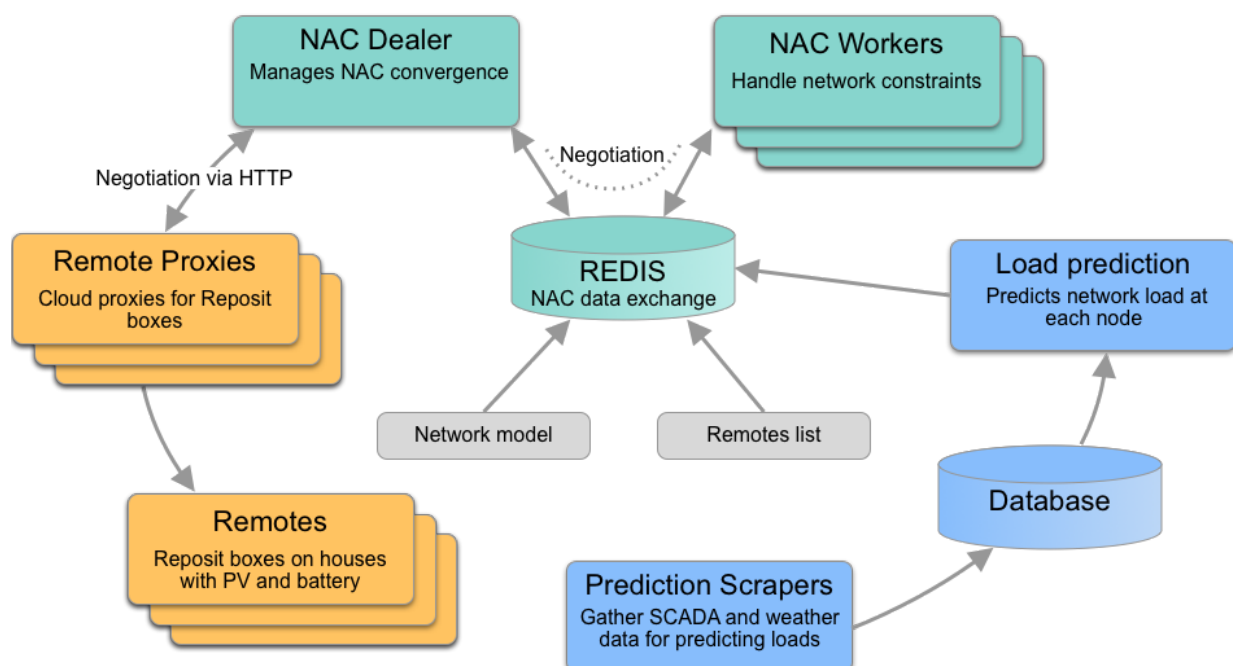






Fig. 4. Detailed NAC architecture, as implemented in CONSORT. Different systems are highlighted in different colours: teal for the NAC, orange for the remotes, and blue for the load prediction.

The NAC dealer, REDIS server, workers, scrapers, database and load prediction services all run on a single server in a data center. The remotes are physical Reposit boxes installed in the participants' houses on Bruny Island; they interact with Reposit's cloud infrastructure, which includes the remote cloud proxy service as well as other services that are not directly visible to the CONSORT project.

Recall that NAC is essentially a negotiation between remotes, which control the dispatch of DER (household batteries), and workers, which ensure that the network constraints are satisfied. The negotiation is mediated by the "NAC dealer". Below, we will discuss implementation details of these elements in turn.

### 3.2.2 Reposit boxes (remotes)

Without NAC, each Reposit box plans the dispatch of its associated battery such that the house's electricity bill is minimised. Often, the battery is charged whenever excess solar power is produced, and discharged to meet household demand during peak periods, when tariffs are higher. The Reposit optimisation occurs according to Reposit Power's own algorithms, which are not directly accessible to the NAC.

As explained in the "NAC Prices" section, above, NAC modifies this calculation by providing additional NAC prices as input to the optimisation. The NAC prices offset the normal tariffs that the household would experience, and act as an incentive to either dispatch or consume extra power<sup>2</sup>.

#### Proxy remotes

Ideally, each Reposit box would communicate directly with the NAC. NAC requires an iterated negotiation to take place, with an associated communications overhead. Because of possibly unreliable, slow or limited internet connections at participant households - a particular problem on semi-rural Bruny Island - a design decision was taken to carry out this negotiation with special "proxy" or "virtual" Reposit boxes running in the cloud. These are designed to act just as the corresponding real Reposit box would. At the end of the negotiation, the final NAC prices are passed to the real Reposit box, which then acts as it sees fit based on these prices.

The use of proxy remotes is simply an implementation detail - for most purposes, the proxy boxes can be considered to be perfect substitutes for their corresponding real boxes.

---

<sup>2</sup> In CONSORT, all support events were to dispatch, rather than consume power, meaning the NAC price was positive. Negative NAC prices would be typical when trying to reduce high voltages, a scenario that wasn't explored in the CONSORT trials.



### 3.2.3 NAC workers

The other partner in the NAC negotiations are the NAC workers. These are agents that are responsible for ensuring that the network constraints are satisfied at the lowest possible cost. In the CONSORT implementation of NAC, the workers have two means of satisfying these constraints: they can either dispatch diesel from the mobile generator - which incurs a cost due to diesel consumption - or they can encourage dispatch from the remotes (participating households), at a cost given by the latest negotiated NAC prices.

As explained in the section on ADMM, there is in effect one worker agent for each timestep in the horizon. Each worker is only concerned with ensuring that the network constraints are met for that timestep only.

Each optimisation call takes as input the static network model, the NAC prices, and the predicted network loads. This information is drawn from data stored in the REDIS database (see below).

Based on this data, a mathematical optimisation problem is formulated. Implementation details may be found in Scott *et al.* (2019) [2]. Briefly, each network element - e.g., lines, loads, transformers, tap changer etc. - has its own equations which contribute to the cost and constraint functions in the optimisation problem. A full three-phase model of the network is used to improve accuracy on this significantly unbalanced feeder.

At each iteration, the workers are provided with a new set of prices, and in return they give a desired set of dispatch levels for the remotes and the diesel which most efficiently satisfy the network constraints, for the given prices.

The NAC worker code is written in C++, and makes use of the MADOPT and IPOPT optimisation libraries.

### 3.2.4 The NAC dealer

The NAC dealer, written in C++, implements the ADMM algorithm. For each horizon, it iterates a sequence of message/reply pairs until convergence, and, therefore, a near-optimal solution to the problem, is achieved. At each iteration, there is a single message/reply to each proxy remote, and a single message/reply to a worker for each timestep in the horizon. With 48 timesteps in each horizon, and 32 participating households, there will be 32 messages to the remote proxies, and 48 messages to the workers.

Convergence is deemed to have been achieved when the planned dispatch levels agree to within a given tolerance between the workers and remotes, and a similar agreement holds for the prices. When these conditions hold, we know that the solution is approximately *feasible* (all constraints are satisfied) and *locally optimal* (cheapest). After convergence has been





achieved, a final message is sent to the remotes, so that the households (Reposit box) may plan their dispatch based on the agreed negotiated solution.

### 3.2.5 REDIS in-memory database

REDIS is an in-memory database cache that is often used in distributed software. For the NAC, it acts both as a shared memory that may be simultaneously accessed by multiple running processes, and as a message queue that is used for coordination and communication between these processes. As shown in Fig. 4, all communication between the NAC dealer and workers is routed through REDIS.

In contrast, the NAC dealer uses the ubiquitous HTTP protocol to communicate with the remotes. This makes sense, because the remotes (Reposit boxes) are an external participant in NAC, and HTTP is more appropriate for this kind of communication across product boundaries.<sup>3</sup>

### 3.2.6 e-JSON data format and the network model

The NAC workers solve an unbalanced power flow problem based on a mathematical model of the network. Due to a lack of suitable alternatives, we chose to develop our own data format for storing and communicating the network model. We call this format e-JSON. It is based on the commonly used general purpose [JSON format](#).

For CONSORT, we have fleshed out the e-JSON format only as far as is necessary to model the Bruny Island 11 kV network. Thus, the e-JSON specification includes provision for modelling various kinds of lines (overhead / underground / underwater / 3-phase / single-phase spurs etc.), an infeeder model, delta and wye loads, generators, open-delta and distribution transformers, time series data, and specialised components representing the participating households.

Our e-JSON network model was initially populated by converting the [SINCAL](#) network model provided by TasNetworks. Some additional modifications were carried out to trim off external pieces of the network model, to telescope together consecutive lines where possible, and to ensure that the correct electrical phasing was used. Following this procedure, we ended up with the network shown in Fig. 1.

#### Network phasing

Due to the usual issues of visibility of distribution data - for example, phase transpositions not being recorded in computer models - additional work was done early in the project to correct

---

<sup>3</sup> Recently, to aid uniformity and flexibility, we have implemented a REDIS-based communication between the dealer and remotes, with the remote to redis messages handled by an extra “remote communicator”.



the phasing given by the SINICAL model. This involved a visit to Bruny Island, talking to engineers, and working with hand drawn diagrams and other notes.

### Phasing of participant houses

There was also an issue with determining the phasing of participant connections. Where the connection originated at a single-phase 11 kV transformer, there was no ambiguity. However, in some cases, customer phasing was unknown and needed to be randomly assigned. We estimate about 75% of the phase assignments used by NAC and in simulations to be correct. As such, the errors should have had a relatively minor impact on the accuracy with which NAC was able to fulfil network constraints.

### Open-delta voltage regulator

Special consideration was also required for the open-delta voltage regulator, see Fig. 1. A special model needed to be developed both for NAC, and for the simulation.

## 3.2.7 Load prediction

To effectively solve the forward horizon, the NAC workers must be able to predict the load at each 11 kV node in the network. In the CONSORT implementation of NAC, the workers draw these load predictions from the REDIS database, and there is a load prediction service that updates this prediction every five minutes.

The load prediction service starts with an aggregate prediction of the total *background* Bruny consumption (not including diesel and net loads for participating households) for the next 24 hours. The prediction has a 30 minute granularity and is updated every thirty minutes. It takes as input recent SCADA data for the total cable import and diesel, recent participant data, and weather forecasts. The database containing this input data is periodically updated by scraper services that query the SCADA system via Reposit Power's API and the BOM's API.

Initially, a linear regression-based predictor was used for the aggregate prediction, but since 16/07/2018 a new neural network based prediction, developed by CONSORT partners University of Tasmania and TasNetworks, has been used [6].

Every five minutes, the load prediction is augmented using very recent SCADA and participant data to achieve a more accurate very near-term forecast. The resulting aggregate forecast is then disaggregated into loads at each 11 kV distribution transformer. The disaggregation simply multiplies the prediction aggregate load by a different constant factor for real and reactive power at each transformer. The constant factors are calculated based on historical data. The final disaggregated loads are then fed into the REDIS database, where they are picked up by the NAC workers.



## 4 Simulations

Simulations have been used throughout the CONSORT project as a means of doing offline testing of the NAC, for situational awareness (e.g., for inferring voltages), and for answering counterfactual questions regarding peak events (e.g., what would have happened during a peak if NAC hadn't been operational? How much money was saved? etc.).

The simulations are written around the [SmartGridToolbox](#) library [5], developed by the ANU CONSORT partner. Various extensions have then been written on top of SmartGridToolbox for the CONSORT project:

- A parser for the e-JSON format has been written,
- An open delta transformer model has been developed (which has been incorporated back into the base SmartGridToolbox library)
- A “Remote” simulation component type has been developed. This represents one of the participating households, and comprises household load and PV generation sources that obtain their data from time series data files, a battery that may either act according to a historical time series, a Reposit proxy server (or an emulation of one), or a Reposit proxy server operating in the presence of NAC
- Various pieces of analysis software and utilities have been developed
- A “ConsortServer” component has been developed, which interacts with various parts of the NAC software and Reposit proxy servers or emulations thereof, allowing the simulation to act as a “virtual network” for an actual running NAC instance. We shall elaborate on this below
- Scripts have been written to coordinate the SmartGridToolbox-based simulation with the various components of the NAC.

### 4.1 CONSORT simulation architecture

The CONSORT simulation architecture is shown in Fig. 5, below. There is a simulation layer, which is based on a SmartGridToolbox simulation. This simulates the state of the network, including both participant and background loads aggregated to the distribution transformer level, as well as the participating households. Running alongside this are the NAC as well as and the Reposit proxy servers.

The proxy remotes run in a simulation mode, in which they do not communicate with their corresponding real Reposit boxes, and are provided with extra data about the predicted household load and PV generation - taken from Reposit's historical database - as well as tariffs and other static data. They run in Reposit Power's cloud infrastructure. Alternatively, we can use a close emulation of Reposit's proxy servers that can run in the NAC data center, to conserve communications bandwidth and boost performance.

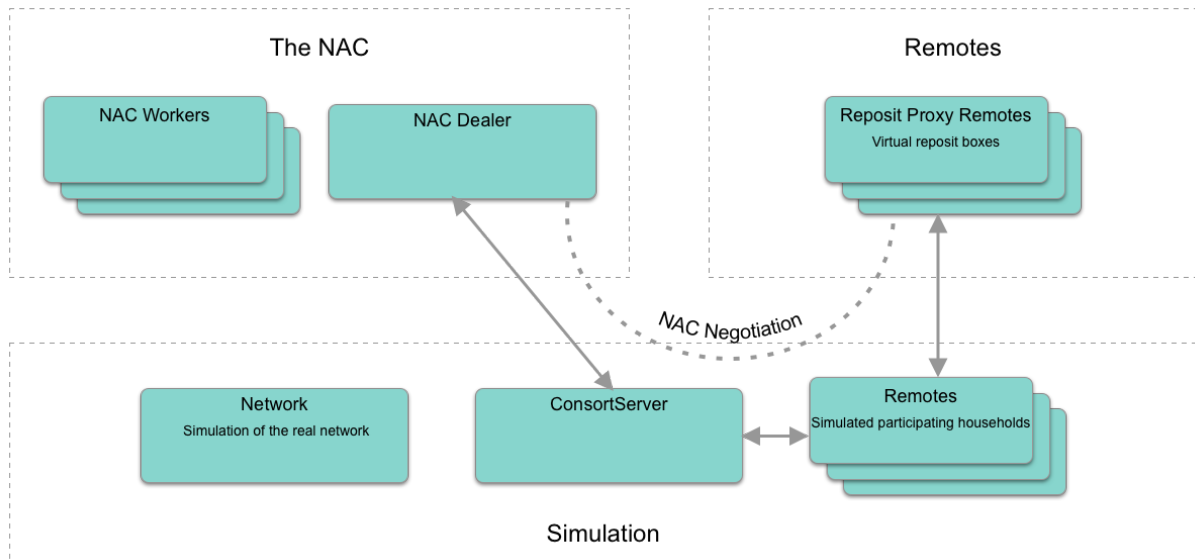


Fig. 5. Schematic of the simulation architecture used in CONSORT.

The NAC is also run in a simulation mode, which is similar to live mode, except that the current time is set according to simulation time, rather than the real time, and, whenever the negotiation corresponding to a horizon finishes, time is “fast forwarded” to the next horizon.

This “fast forwarding” behaviour is echoed in the simulation. During NAC negotiations, the simulation advances in time at a rate of one simulation second per real second. When the simulation receives a message from the NAC dealer indicating that the horizon is finished, the simulation is instead allowed to progress quickly to the next horizon.

## 5 Trials

### 5.1 Goals from NAC perspective

From a perspective of developing the NAC, the main purposes of conducting live trials were the following:

- To demonstrate the viability of NAC for controlling congestion in the network. In essence, we wanted to demonstrate that the NAC was viable for solving Bruny Island’s main network issue
- To gain insights about the behaviour of the NAC algorithm when used in realistic circumstances
- To provide a baseline and validation for simulations designed to answer counterfactual questions like: What would have happened during the trials, if the batteries and PV had



been present but NAC had not been used? What if there were no batteries / PV? What would be the effect of adding extra participants?

- To study the performance and scaling of the NAC algorithms
- To study the use of NAC to control network voltages.

## 5.2 Relevant trial periods

To date, the NAC has been trialled in live operations for a total of 83 days, as shown in Table 1 below. These cover trial periods where peaks are known to have historically occurred [3], such as long weekends in colder periods and school holidays, as well as some special trials that were run to test features of NAC. There were 17 true network peaks, and 30 network support payments. The extra 13 network support payments were mainly due to “artificial peaks” that were created for various testing purposes by either lowering the threshold for which the cable current was deemed to be too high, or by putting in place artificial voltage limits.

## 5.3 Supplementary simulations

Simulations were performed using data collected from some of the periods listed above, to answer counterfactual questions such as: What would have happened during the trials, if the batteries and PV had been present but NAC had not been used? What if there were no batteries / PV? What would be the effect of adding extra participants? The results of these simulations will be given in the “Results” section below, alongside results from the corresponding actual trials.



## 6 Results

### 6.1 Summary of NAC trial periods

Table 1 lists the 8 NAC trial periods over 2018-19, along with information about what occurred.

Code <sup>4</sup>	Name	First Day	Last Day	Days	Peaks <sup>5</sup>	Payments <sup>6</sup>	Notes
TP180329	Easter	29/03/18	03/04/18	6	6	7 x ER	
TP180413	Apr. sch. hols	13/04/18	03/05/18	21	0	0	Test NAC for extended period.
TP180608	Queen's B'day long w'end	08/06/18	12/06/18	5	5	5 x ER	
TP180706	July sch. hols	13/07/06 <sup>7</sup>	23/07/23	11	5	2 x ER, 9 x EU	Artificially low line limit for much of period - see results.
TP181219	Summer hols	19/12/18	09/01/19	22	0	0	Test NAC for extended period; very conservative voltage limits (not expected to be triggered)
TP190125	Australia Day	25/01/19	29/01/19	5	1	1 x EU	Conservative voltage limits
TP190130	Special vlims	30/01/19	08/02/19	10	0	3 x EU	Less conservative voltage limits leading to network support payments
TP190208	Special SV	08/02/19	09/02/19	2	0	3 x EU	Special run with artificially low line limit in NAC to test Shapley-value based payments

Table 1. Summary of NAC trials.

We summarise in more detail the results for each of these periods in the following sections.

<sup>4</sup> TPYYMMDD where YY = year, MM = month, DD = day.

<sup>5</sup> # of network peaks that would nominally require diesel dispatch in the absence of NAC.

<sup>6</sup> # of Network support events involving payments to customers. ER = energy reserve, EU = energy use. May include special payments e.g. artificial peaks for testing purposes, special voltage peaks.

<sup>7</sup> From 6/7 to 13/7 NAC was run, but due to an issue, Reposit boxes ignored NAC. This period has therefore been excised from the table.





## 6.2 Details of trial periods

### 6.2.1 Easter 2018 (TP180329)

This was the first full-scale trial of NAC carried out on Bruny Island. A graph summarising what occurred is shown below.

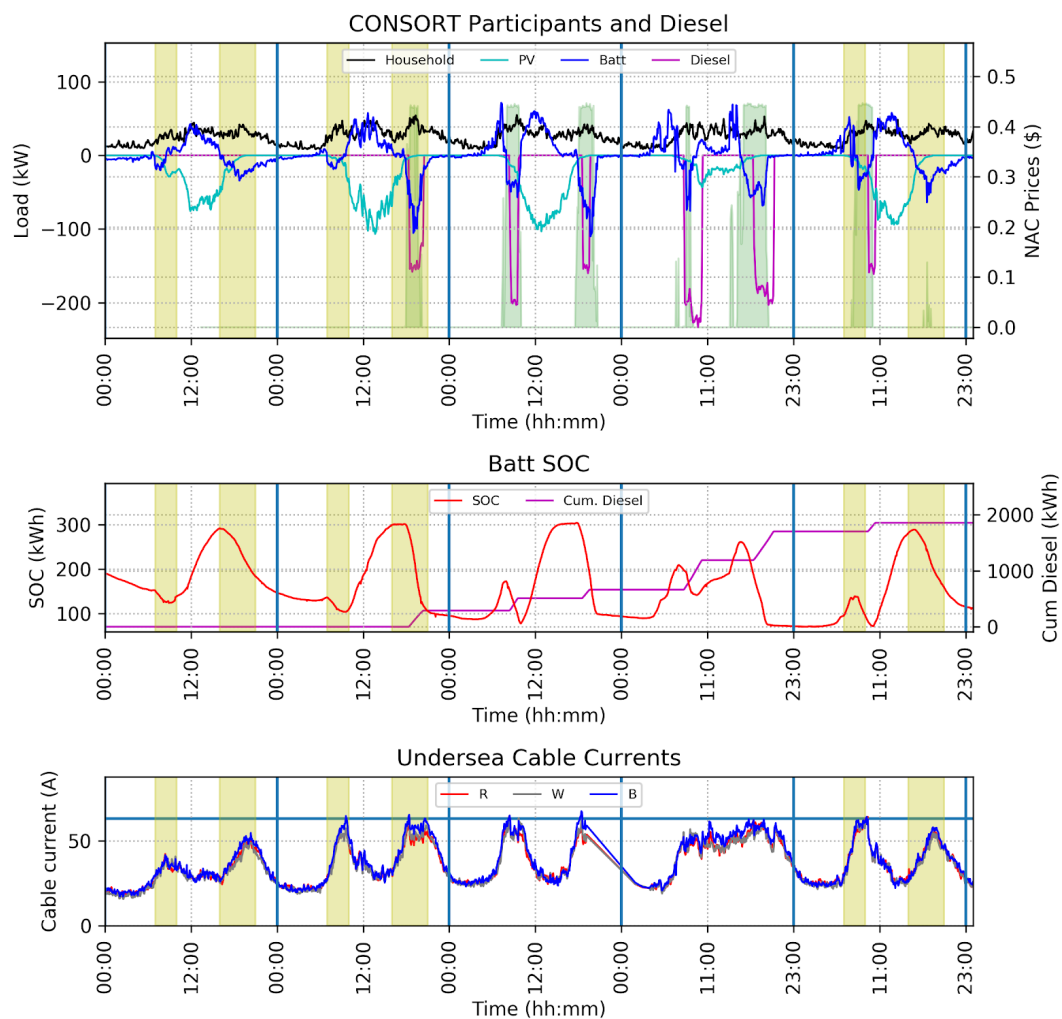


Fig. 6 Easter Trial, TP180329. Top: aggregate participant loads, diesel load; negative = generation. Yellow bars are peak TOU tariffs. The pale green filled graph represents non-zero NAC prices. Middle: aggregate battery state of charge, cumulative diesel discharge. Bottom: current in the undersea cable. The horizontal line is the approximate upper limit to which the current is normally controlled by the diesel operators (63 A).

The top graph shows the aggregate loads for all participants, as well as total diesel generation, in the trial over the course of five days, starting on 29 March 2018. The vertical



light yellow bars represent periods when TOU peak tariffs were in place, and therefore when it becomes most advantageous for participants' batteries to discharge to meet household load. There is no peak pricing on weekends, days 3-4 in the figure. Participant loads are split into household, PV and battery components. Negative loads (for PV and sometimes battery) represent generation. On days 2-5, a light-green shaded plot can be seen. This is the NAC prices, and corresponds to the right hand axis of the graph. Positive NAC prices provide extra incentive for batteries to discharge.

The middle graph shows the total state of charge of all the batteries (red, left hand axis) and the total diesel discharged during the period in question (magenta, right hand axis).

The lower graph shows the line currents on each phase through the undersea cable. The horizontal line, at 63 A, is the approximate current to which the control room operators try to limit the current when dispatching diesel during peak events.

### 6.2.2 April school holidays (TP180413)

This extended test period of 21 days didn't see any network peaks, but allowed the NAC to be continuously tested over a long period, thus providing reassurance against the kinds of bugs that may only manifest after a long time (e.g. memory leaks and overflows) or on rare occasions. Fortunately, no issues were observed in the operation of the NAC, and the test passed without any problems.

### 6.2.3 Queen's Birthday long weekend (TP180608)

The Queen's Birthday long weekend is typically another period where major network peaks are seen, due to it being a long weekend that occurs during the cold seasons. A graph showing the trial period is shown below.



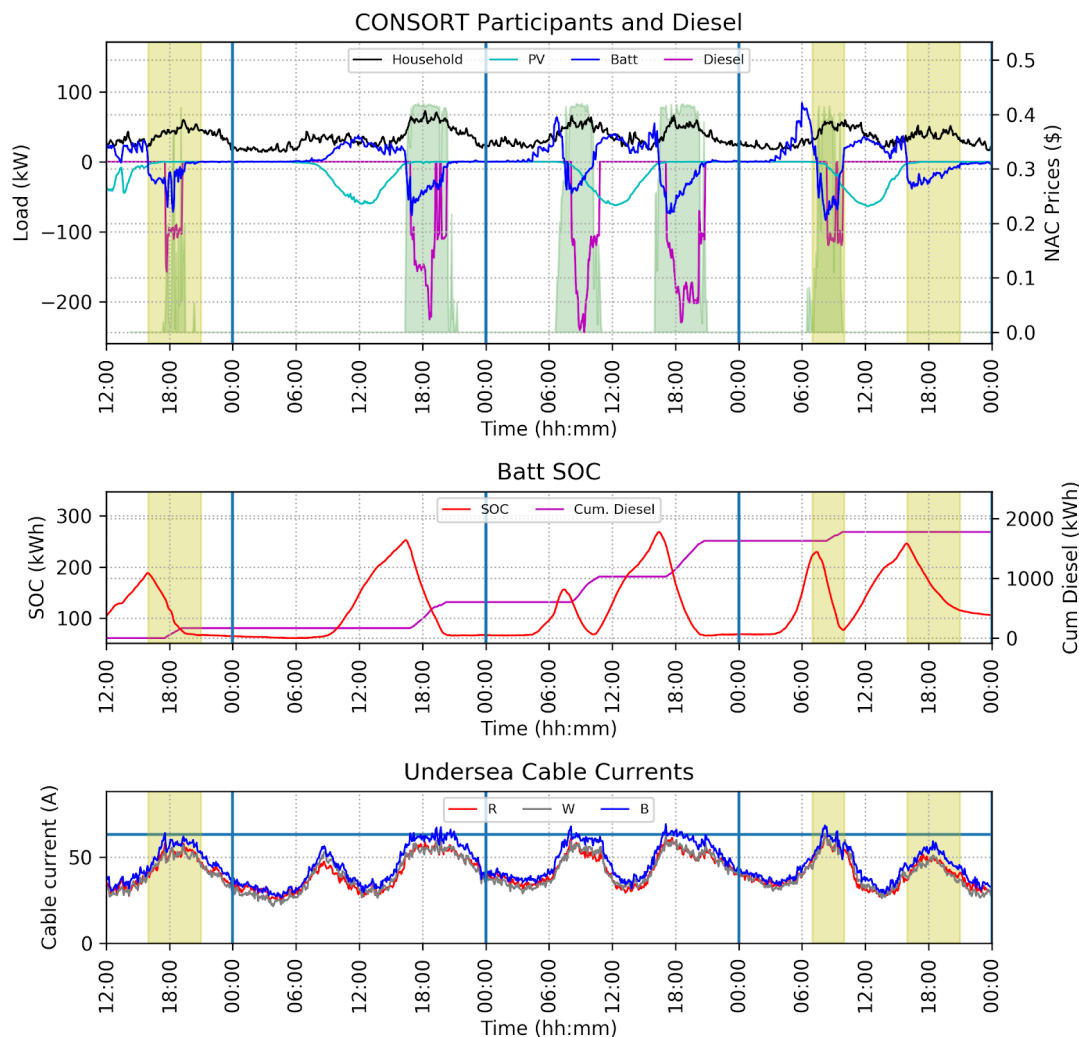


Fig. 7. Queen's Birthday long weekend trial.

Prior to this trial, some improvements in NAC were implemented: an improved set of parameters for the disaggregation of load predictions, and an improvement on the way the very short-term load forecasting was carried out.

Compared to the Easter 2018 trial period, we can see that the cable current is limited more tightly to the 63 A upper bound. This is probably due to the improvements mentioned above as well as differences in the way diesel was dispatched.

We also saw that the batteries were capacity limited during the evening peaks, and that the batteries did not charge enough overnight during the morning peak of the third day. We believe this is due to the peak not being predicted far enough ahead of time, possibly interacting with settings that caused the Reposit controllers to act as late as possible, a feature that is deemed desirable under normal (non-peak) circumstances, but which may



occasionally reduce the effectiveness of NAC. Some tweaking of these settings could potentially prevent this occurring in the future, as could having better forward prediction of peaks.

#### 6.2.4 July school holidays (TP180713)

During the July 2018 school holiday period, we experimented with lowering the cable limit in NAC, to artificially force it to dispatch the batteries. The main aim was to do back-to-back trials of the two main payment methods used in CONSORT: the *energy reserve* method (payment calculated prior to the peak based on expectation; customers are paid for reserving energy), and the *energy use* method, (payment calculated after the peak, and customers are paid according to actual energy dispatched for network support).

In addition, we took the opportunity on the 16th July to switch to a new neural network-based forecasting method [6].

In the period between 6 July and 13 July, it became apparent that the Reposit boxes had not correctly switched into NAC mode and were effectively ignoring all NAC message. This period does not therefore represent a valid test of NAC, and has been excluded from analyses of the NAC's effectiveness.

Fig. 8, below, shows a highlight of the July holiday period, when NAC was able to unambiguously prevent a diesel start for the first time.

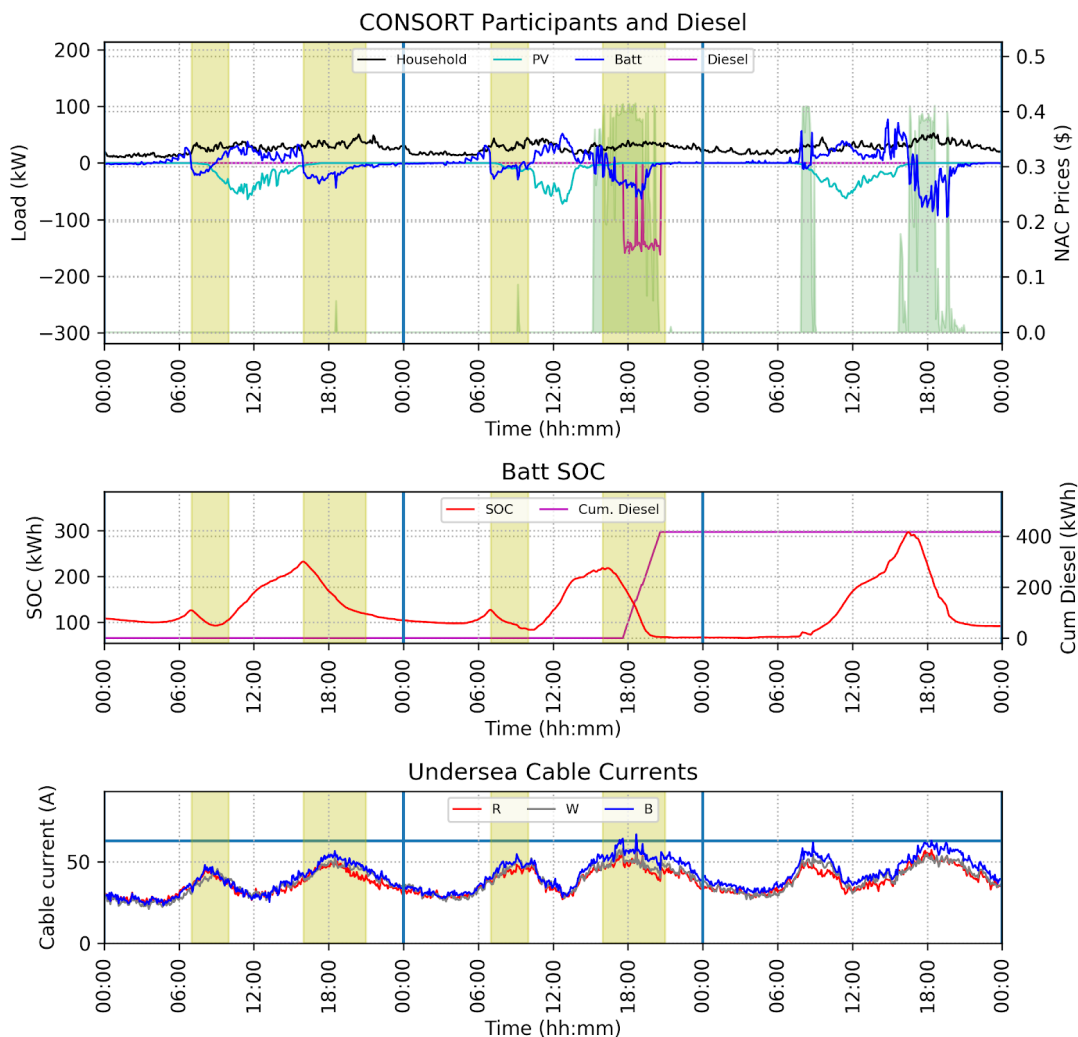


Fig. 8. July 19-21 2018. A diesel start was prevented during the evening peak of the third day.

### 6.2.5 Summer holidays 2018-19 (TP20181219)

This period served as yet another extended test of NAC, with some hope that it would be able to be of use during a peak or two. In the end, there were no network peaks during this period.

In addition, the use of lower voltage limits was trialled for the first time. NAC had always been equipped to be able to handle network voltages, but prior to this point, controlling the current through the cable provided the focus, and voltage limits were not used. For this test period, we placed voltage limits of 0.96 p.u. near two nodes at the southwest and southeast extremities of the 11 kV network (see Fig. 9 below). This value was very conservative, in the sense that voltages were not expected to drop this low unless there was a major network peak event.

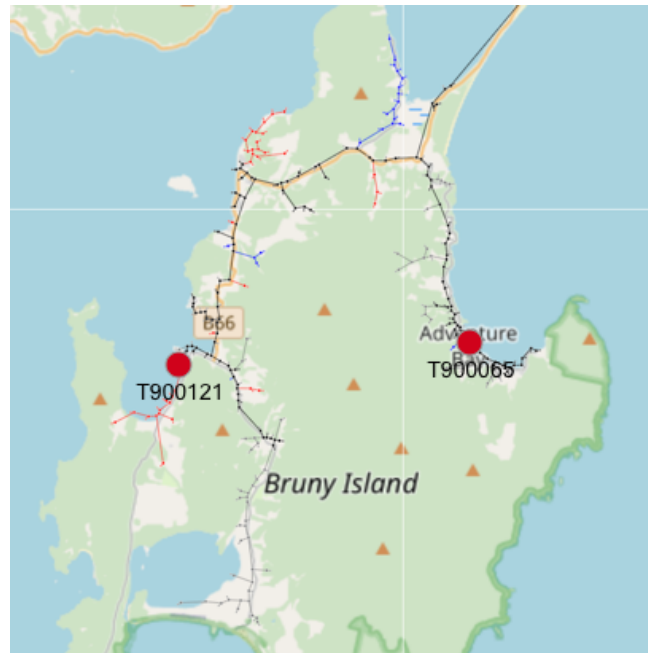


Fig. 9. Location of voltage controlled nodes

### 6.2.6 Australia Day (TP190125)

During this period, settings were the same as TP20181219. A single network peak occurred, during which The NAC performed as expected.

### 6.2.7 Voltage control tests (TP190130)

In early 2019, the voltage lower bounds used in the previous two tests were raised to 0.99 p.u. The results will be discussed in a separate subsection below.

### 6.2.8 Shapley value payments (TP190208)

For the final test, an artificially low line limit of 40 A, to perform a dry-run of the scripts used to calculate, process and send Shapley-value based payments and associated messages to participants. These scripts were run in parallel to the usual “energy-use” based payment scheme. Live Shapley value calculations were performed during the trial, and the results collated for sending to participants, but the actual messages and payments were generated based on energy use calculations. The Shapley-value payments were found to pass a sanity check, in that they confirmed expectations of what should happen based on the details of batteries, phasing, network locations etc.

## 6.3 Simulations and further analysis

A detailed analysis was carried out, involving the Queen's Birthday long weekend (TP181608, 8-11 June) and the July school holidays (TP180713, 21-22 July). First, to provide a fair baseline for counterfactual simulations, and to validate the accuracy of the simulation, a complete replay of each trial period in the analysis was performed. As described earlier, a SmartGridToolbox-based simulation layer was used to simulate the network, including participant household loads, PV and batteries. The NAC and proxy remote solvers were run alongside this simulation, with the simulation communicating with these to determine the power dispatched from batteries, according to the joint solution negotiated by the NAC and proxy remote solvers. This process is able to provide an environment that is quite close to what was actually used during live operation of the NAC.

An example of one such baseline simulation, for the Queen's Birthday long weekend, is shown below.

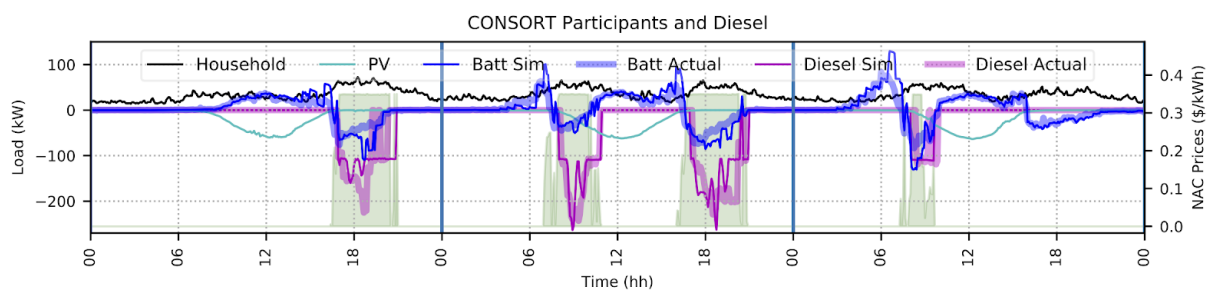


Fig. 10. Queen's Birthday long weekend, June 9-11, comparing actual and simulated values. Actual values are shown as thick, light lines, and simulated values are shown as thinner, darker lines. The simulated diesel is calculated using a simulation model of human-in-the-loop control. Source: P. Scott *et al.* (2019) [2].

For the most part, the battery actions follow the historical recordings fairly closely. Disagreement could be due to differences in the way the historical predictions to the batteries are handled, the simulation having to fill in "gaps" in the prediction and historical data due to data losses (e.g. because of households having been out of internet communication at the time of the trial), or other issues around internet connectivity, or timing differences in the way the NAC negotiations played out.

Having thus generated a baseline simulation, further counterfactual simulations were then carried out. The complete list of cases for each period in the analysis is given below.

1. Actual, historical trial recordings (NAC)
2. Simulation replaying NAC (NAC Sim)
3. NAC with perfect island load forecast (NAC Perf)
4. Manually scheduled dispatch strategy (Manual)



5. Independent local battery optimisation (Indep)
6. Participants without batteries (No Batts)
7. NAC with 100 batteries (NAC 100).

The “NAC Perf” case allows us to estimate the maximum benefit that could be derived by improving the forecasting module. The “Indep” case uses a scheduling strategy that mimics manually locking in a battery discharge prior to an anticipated peak; many such pre-programmed dispatches were performed during the project, using the Reposit Fleet application. The “Indep” case allows us to assess the value that the NAC has added on top of simply having the PV and batteries in the network, with no special coordination. The “No Batts” case allows us to explore the value of the batteries plus NAC, and the “NAC 100” case gives an idea of how effective increasing the number of participants might be.

Fig. 11, below, compares the battery and diesel discharge for four of these scenarios, for a single day.

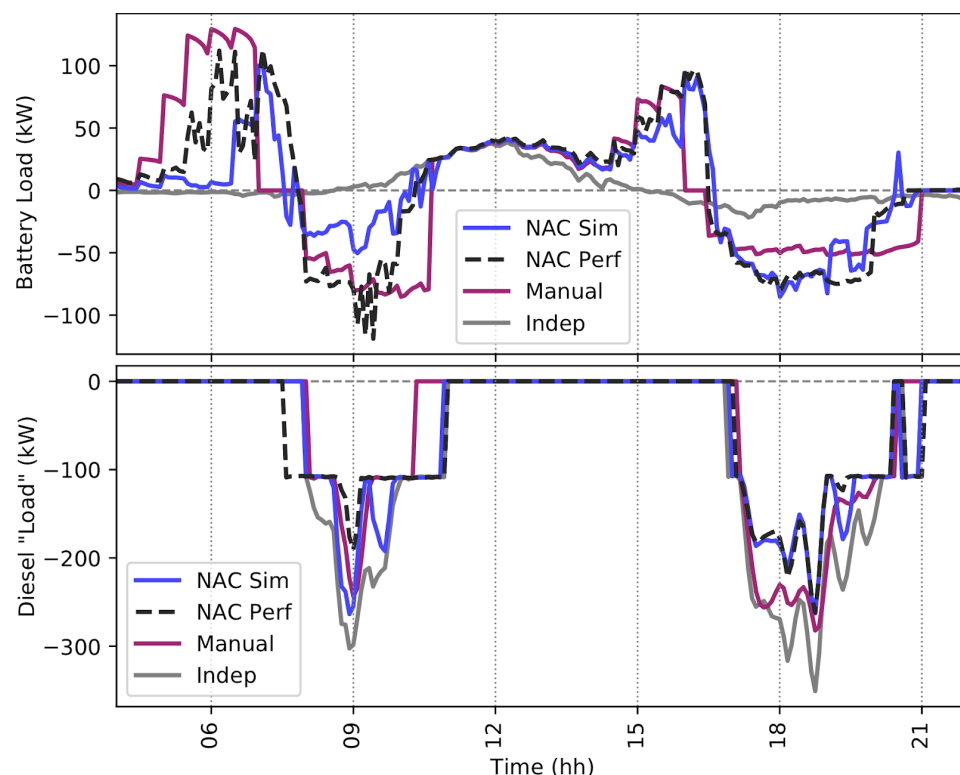


Fig. 11. Battery load (top) and diesel “load” (bottom) for June 10 2018. Source: P. Scott *et al.* [2].

Of particular note, the “Indep” battery discharge is much smaller, in terms of discharge, than all the others for both peaks; the “NAC Sim” underperforms compared to the “NAC Perf” during the morning peak, due to not predicting the peak in time to fully pre-charge; and the “Manual” battery discharge is too great during the tail of the morning peak and not great enough during the evening peak, demonstrating clearly the value of the NAC.



Fig. 12, below, shows a cost analysis of these simulations (plus the recorded historical values). The analysis was done on the combined results of Queen's Birthday 2018 results (June 8-11) and July school holidays (July 20-21).

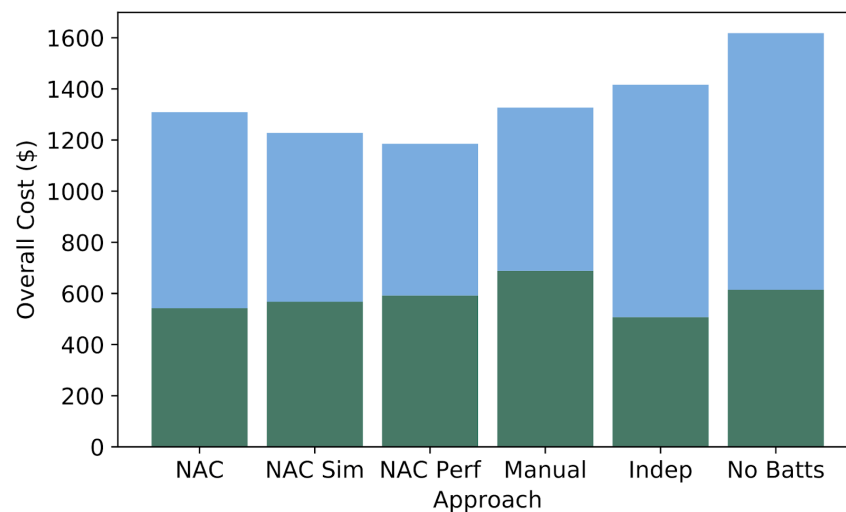


Fig. 12. Cost analysis for Queen's Birthday long weekend (June 8-11) and July school holidays (July 20-21). Green = participant costs (mainly tariffs, and not including network support payments) and blue = diesel costs. Source: P. Scott *et al.* (2019) [2].

The actual trial (NAC) has a 7% higher cost than the simulation (NAC Sim). This is because the diesel generator was left on for longer than strictly necessary by the TasNetworks operator during the July peak. After taking this into account, the simulation agrees with the live trial total costs to within 0.7%.

For the two trial periods analysed here, the methodology of counterfactual simulations described above suggests that the NAC achieves a 13% reduction in total costs compared to uncoordinated batteries ("Indep"). The participants incur a slightly higher cost so that the diesel can be significantly reduced. No matter what payment scheme is chosen, care should be taken for the network to adequately compensate customers in such cases.

The NAC achieves a 7.5% improvement over the Manual schedule.

The NAC Perf case represents the response that could be achieved with perfect non-participant load forecasts. It indicates there is room for a modest 3.5% improvement if a better forecasting engine could be implemented.

The 31-32 batteries participating in NAC during the period of analysis, which represent around 4% of all feeder customers, have the capacity to meet 10% of peak feeder demand. Through the use of NAC, they managed to achieve a 34% reduction in diesel use compared to the case with no batteries, for the periods in question. To see how this changes with the number of



batteries, a NAC simulation with 100 batteries was conducted, which reduced the diesel consumption by 74%.

## 6.4 Voltage control

As mentioned above, the voltage at two nodes (see Fig. 9, above) was controlled by NAC during trial periods TP181219, TP190125 and TP190130. The limits were set conservatively during the first two periods, such that a response was not expected. During the third period, TP190130, the limit was tightened from 0.96 p.u. to 0.99 p.u., which was expected to induce some artificial network support events.

Voltage limits were “soft” limits, meaning that the NAC was allowed to violate them, but at a “cost” comparable to but below the cost of diesel. As such, there would be a strong incentive to dispatch batteries to control network voltages, but not so strong that the NAC would attempt to discharge diesel<sup>8</sup> to meet these constraints.

The following image shows the operation of NAC on 2-5 Feb, TP190130. Two minor battery discharges, coupled with non-zero NAC prices, can be seen during the evening peak of the final day: we have therefore observed the batteries responding to a voltage control event during live trials.

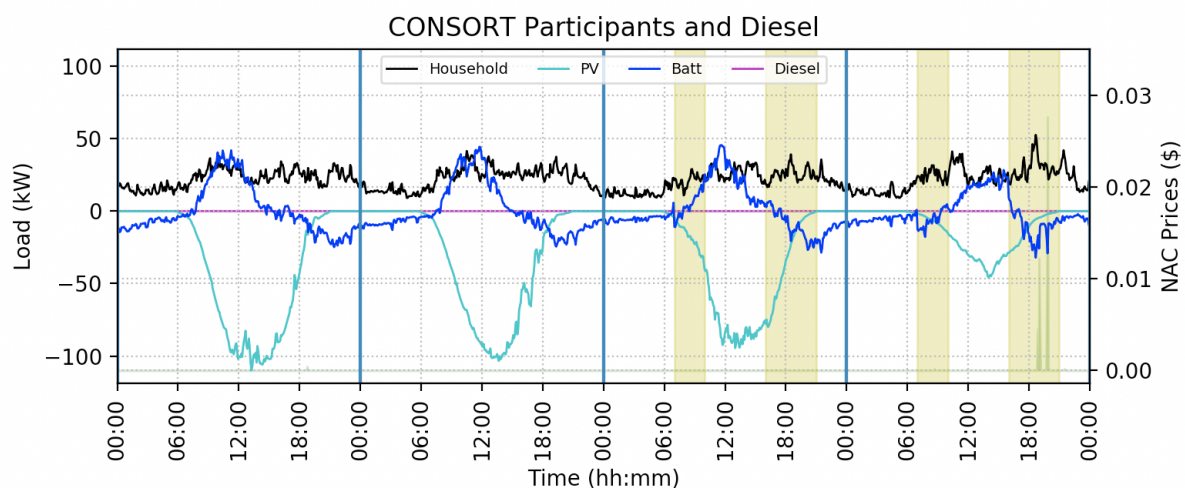


Fig. 13. NAC aggregate powers, TP190130, 5 February.

It is not straightforward to correlate these minor voltage support events with changes observed in the network voltages. One reason for this is the lack of visibility over voltages in the Bruny distribution network. The 11kV voltages are monitored on two of three phases at the tap

<sup>8</sup> The NAC doesn't actually discharge diesel in CONSORT; this is wholly under the control of the control room operators. However, it does keep track of virtual or expected diesel discharge as part of its network modelling.





changer, relatively near to the undersea cable, at a sectionalizer in the northwest of the island, and two reclosers immediately adjacent to each other near the top of the southern part of the network. The sectionalizer and recloser voltages are known to be unreliable - for example, the two reclosers exhibited a constant voltage offset of 0.2 kV relative to one another, despite being immediately adjacent in the network.

There is another issue with modelling network voltages: the recordings show an anomalously high WB voltage, which seems to be consistent for all readings below the open-delta voltage regulator, and which disappeared at one point when the voltage regulator went out of action. This suggests that there is some issue with the voltage regulator. Trial and error simulations show that we can approximately reproduce the issue by placing a large series reactance on one of the two regulator tanks. However, what this all means remains uncertain.

Because of limitations with the voltage recordings etc, it is not possible to directly observe the effect of the discharge shown in Fig. 13 on network voltages. Instead, we have performed a series of simulations to quantify the effect of this discharge, as well as to explore what would have happened under less conservative voltage control settings.

Below, we show plots at one of the two voltage controlled nodes, T900065, on 8 Feb 2019, which is the final day shown in Fig. 13.

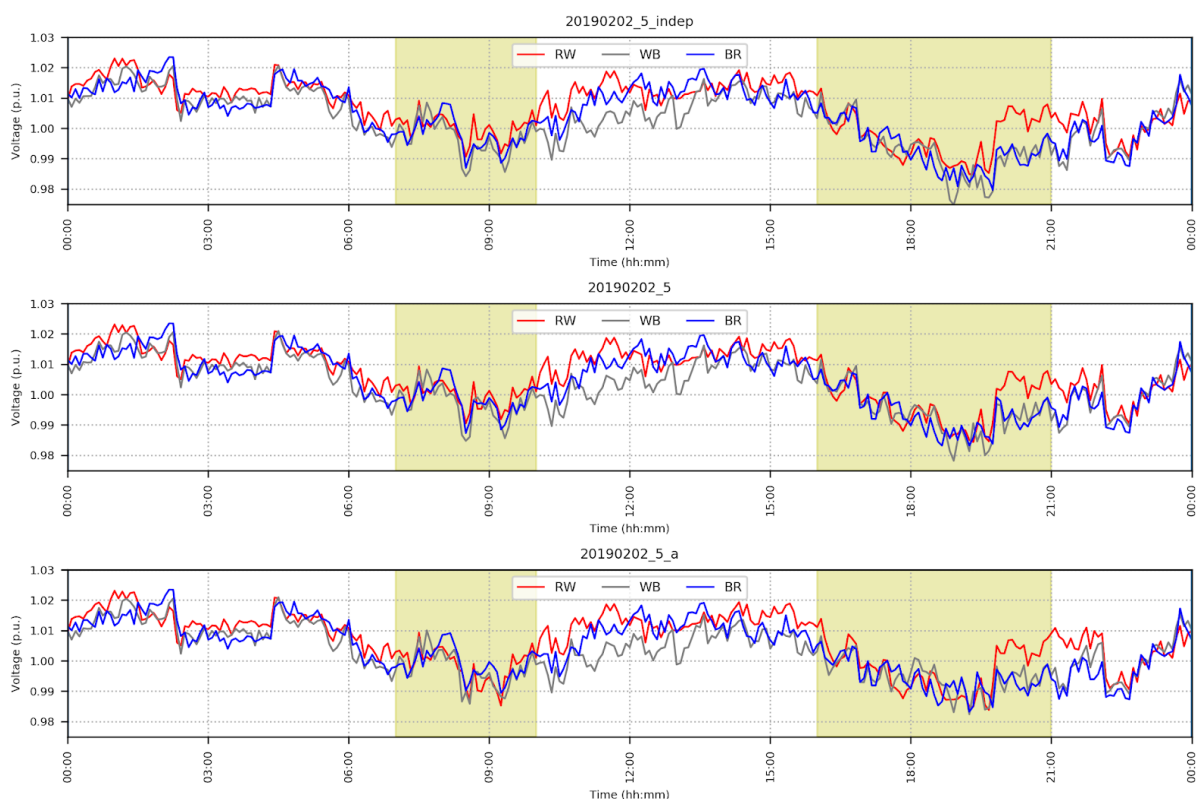




Fig. 14. Simulated voltage at T900065 (See Fig. 9), 8 Feb 2019. Top: no voltage control, middle: voltage LB = 0.990 p.u., bottom: voltage LB = 0.995 p.u.

The top graph shows the simulation results if NAC is not used and the batteries act independently. The middle graph shows the results for the voltage limits used in the real trial, 0.99 p.u. The bottom graph shows the result if the voltage limits are raised to 0.995 p.u. A subtle raising of the lower voltage floor is apparent during the evening peak as we move from the top graph to the bottom. Thus, NAC does indeed appear to be successfully at controlling the voltages, although further work is required before NAC can be used regularly for practical voltage control.

Importantly, voltage control should ultimately be targeted at the low voltage part of the network, rather than the medium voltage network, as was done here. A single battery acting near the end of a low voltage feeder could potentially have a much greater effect on the local voltages than it would have upstream of the distribution transformer. In the CONSORT project, there was little visibility into the low voltage network. There is scope for further work to explore the potential for NAC to directly control LV voltages.

## 6.5 Performance and computational burden of NAC

During the Queen's Birthday trial, TP180608, the NAC converged to within the desired tolerance in 960 out of 982 horizons, with the remaining 22 also solved to acceptable values<sup>9</sup>. We can estimate the "error" in total dispatched battery power corresponding to the convergence parameters. In the worst case, this was found to be 80 W total for all customers, and on average less than 1 W. Consequently, in the future, it would be reasonable to reduce the stringency of the tolerances and thereby reduce computational intensity.

Over all horizons in the Queen's Birthday trial, the mean number of iterations was 18.7, the standard deviation was 15, and the maximum was 64. The average time per iteration was approximately 4 seconds.

The average solve time of each worker job was 316 ms. With 12 workers running in parallel, and 48 problems to solve per iteration, this accounts for around 1.3 s of the 4 s iteration time, while the other 2.7 s is taken up by the dealer and remote computations and communications.

Under normal (no NAC) operation, each Reposit box (remote) re-optimises the forward horizon once per minute. The remote proxies, used during deployment of NAC, perform a similar computation for each iteration in a horizon. If there is one horizon every five minutes, and there are on average  $n$  iterations per horizon, the computational task of the remote proxies is therefore a factor of  $n/5$  the size of that of the reposit boxes themselves. For the

---

<sup>9</sup> "Acceptable" convergence means the horizon ended before the desired tolerance was reached, but the best solution up to that point was good enough to be accepted.



Queen's Birthday trial period, this is a factor of around 3.75 increase in the total computational load of the remote plus remote proxy infrastructure.

During times when the network was not under stress, the NAC was observed to converge within about 2 or 3 iterations per horizon, and the computational burden is greatly reduced.

In the 100 battery simulations, the mean number of iterations was actually slightly better, at 13.0. The mean worker solve time was slightly worse, at 357 ms per job, but much less than the threefold increase in the number of participating customers. Worker tests with 250 and 500 battery instances have produced mean solve times of 390 ms and 430 ms. Although further work is needed, these tests are positive with regard to the ability of NAC to scale to many more customers.

We have estimated the total extra computational and communications costs required to run the NAC to be \$4.5 per day during times of network stress.

## 7 Conclusion

The CONSORT project has brought NAC from a stage of basic research, to a design and robust implementation that runs in live network operations. Through this we have demonstrated that it is not only possible to solve a complex distributed optimisation problem in the real world, but also that it achieves better outcomes compared to available alternatives for using DER to manage network constraints.

We can quantify the contribution of the NAC as follows: during the Bruny Island Battery Trial, we have demonstrated a 13% reduction in combined diesel and participant operating expenses compared to uncoordinated DER reacting simply to retail time of use tariffs. Compared to manually scheduled DER operation, there is a 7.5% reduction in costs, which is in addition to fully automating the process and freeing up operator time for other tasks. Compared to the case if there were no batteries, during the latter half of 2018 the NAC was able to achieve a reduction of around 34% in the amount of diesel consumed.

Looking beyond the Bruny Island Battery Trial, automated techniques that optimally utilise DER, at least cost to networks and least inconvenience to customers, will be critical to ensure our networks can incorporate large amounts of distributed rooftop solar and electric vehicles. We have demonstrated the core technology that will allow the NAC to fulfill this role, but there is still further research required to fully understand its performance in other settings, in particular in regards to scaling to larger numbers of participants and networks, and understanding the overall integration with other parts of the power system including the wholesale market.



## 7.1 Recommendations going forward

### 7.1.1 Voltage control

Controlling network voltages will be an important task for NAC in the future. This requires good visibility of network data and ideally some control over inverter reactive power dispatched by inverters.

In demonstrating NAC for the first time, CONSORT was primarily concerned with a somewhat easier task of handling congestion in a single undersea cable. While we demonstrated some voltage management toward the end of the project, we were limited to working with real power, and the limited number of batteries on the island and presence of a network-level voltage regulator made it difficult to show a strong influence on voltage.

Future work should look into enabling the reactive power capabilities of the inverters for an even stronger control over voltages, improving awareness of network voltages, either by direct measurement or inference, and improving convergence of the NAC algorithm in the presence of voltage constraints.

### 7.1.4 Tools for discovering network state and configuration

It is well-known that distribution networks suffer from a visibility problem. The data about network configuration is often poor, patchy, or only exists in a variety of formats, including hardcopy. Phasing is often unknown. Network metering is patchy, and may be of low accuracy. It may not always be possible to directly discover the complete switching state of the network.

To a large extent, these are problems that may take a long time to be resolved, so solutions such as NAC will simply have to deal with the situation as it exists. There is an opportunity to develop tools and techniques that make the most of available information. These tools could combine information from households (smart metering / EMSs etc.) with SCADA data and known network topology, to fill in the gaps as far as possible. Alternatively, new models for representing networks in the presence of such uncertainty could be developed.

### 7.1.2 Scaling

In future settings, both the number of participants and size of the network are likely to increase beyond what was demonstrated on Bruny Island. While it is possible that the NAC approach, with some changes, could directly handle these increases in size, it is worth also investigating how well-designed approximations could reduce the complexity without significantly degrading performance.

The NAC has been designed so that it can work with both low and high fidelity models. One question that needs to be addressed is to what degree the LV network needs to be modelled



in order to successfully coordinate DER. Expressed another way, what is the impact on costs or constraint violations for successively more approximate LV network models? A similar question can also be posed for the MV network model. In particular, when is it worth the effort (in terms of obtaining / managing data and the extra computational resources) to do unbalanced 3-phase network modelling? With access to the right data and a suitable set of scenarios for DER deployment, it should be possible to address these questions through simulations involving the NAC.

### 7.1.3 Market integration

The NAC solution on Bruny Island minimises the combined cost of participant retail tariffs and diesel use, while meeting network constraints. DER also presents a massive opportunity to improve the efficiency of the wholesale markets. In order to do so, we need to either make the DER an active participant in the wholesale markets, either through their retailer or aggregator, or at least price responsive to the wholesale spot market.

In particular, it will be important to integrate the NAC solution with AEMO's NEM dispatch engine (NEMDE), either through its existing bid and market clearing structures, or to explore modifications to the NEMDE and NAC to enable a joint optimisation of the wider generator dispatch and network stability issues at the transmission level, with the DER capabilities accounting for distribution network constraints.

## References

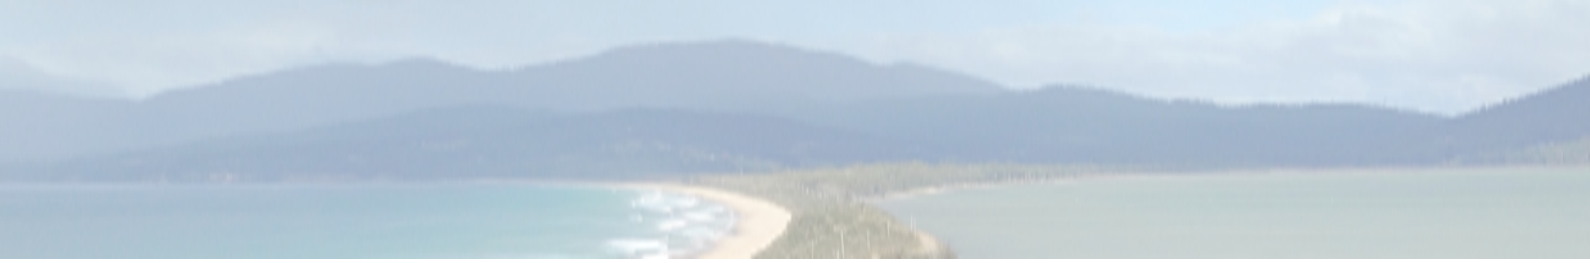
[1] <http://Repositpower.com>

[2] P. Scott, D. Gordon, E. Franklin, L. Jones, and S. Thiébaux, Network-Aware Coordination of Distributed Energy Resources, IEEE Trans. Smart Grid (accepted, in press: <http://doi.org/10.1109/TSG.2019.2907128>).

[3] E. Franklin, D. Gordon, D. Jones, P. Scott, L. Blackhall, and S. Thiebaux, "Peak demand management on distribution networks using coordinated behind-the-meter pv / battery systems: The bruny island battery trial," in 4rd Asia Pacific Solar Research Conference. APVI, 2016, pp. 1–10.

[4] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," Foundations and Trends in Machine Learning, vol. 3, no. 1, pp. 1–122, Jan 2011. [Online]. Available: <http://dx.doi.org/10.1561/22000000016>

[5] SmartGridToolbox (Software Library). Available: <http://gitlab.com/SmartGridToolbox/SmartGridToolbox>



[6] M. Jurasovic, E. Franklin, M. Negnevitsky, and P. Scott, "Day Ahead Load Forecasting for the Modern Distribution Network – A Tasmanian Case Study," in The Australasian Universities Power Engineering Conference (AUPEC), Auckland, New Zealand, 2018.

[7] A. Chapman, S. Mhanna, and G. Verbič. [Reward Structures](#), Final Report, CONSORT Bruny Island Battery Trial Project. April 2019.